Optimal Trade-offs in Multi-Processor Approximate Message Passing

Junan Zhu, *Student Member, IEEE*, Dror Baron, *Senior Member, IEEE*, and Ahmad Beirami, *Member, IEEE*

arXiv:1601.03790v2 [cs.IT] 7 Nov 2016

***Abstract*—We consider large-scale linear inverse problems in Bayesian settings. We follow a recent line of work that applies the approximate message passing (AMP) framework to multi-processor (MP) computational systems, where each processor node stores and processes a subset of rows of the measurement matrix along with corresponding measurements. In each MP-AMP iteration, nodes of the MP system and its fusion center exchange lossily compressed messages pertaining to their estimates of the input. In this setup, we derive the optimal per-iteration coding rates using dynamic programming. We analyze the excess mean squared error (EMSE) beyond the minimum mean squared error (MMSE), and prove that, in the limit of low EMSE, the optimal coding rates increase approximately linearly per iteration. Additionally, we obtain that the combined cost of computation and communication scales with the desired estimation quality according to $O(\log^2(1/\textbf{EMSE}))$. Finally, we study trade-offs between the physical costs of the estimation process including computation time, communication loads, and the estimation quality as a multi-objective optimization problem, and characterize the properties of the Pareto optimal surfaces.**

***Index Terms*—Approximate message passing, convex optimization, distributed linear systems, dynamic programming, multi-objective optimization, rate-distortion theory.**

I. Introduction

A. Motivation

Many scientific and engineering problems [3, 4] can be approximated as linear systems of the form

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{z}, \tag{1}$$

where $\mathbf{x} \in \mathbb{R}^N$ is the unknown input signal, $\mathbf{A} \in \mathbb{R}^{M\times N}$ is the matrix that characterizes the linear system, and $\mathbf{z} \in \mathbb{R}^M$ is measurement noise. The goal is to estimate $\mathbf{x}$ from the noisy measurements $\mathbf{y}$ given $\mathbf{A}$ and statistical information about $\mathbf{z}$; this is a *linear inverse problem*. Alternately, one could view the estimation of $\mathbf{x}$ as fitting or learning a linear model for the data comprised of $\mathbf{y}$ and $\mathbf{A}$.

When $M \ll N$, the setup (1) is known as compressed sensing (CS) [3, 4]; by posing a sparsity or compressibility requirement on the signal, it is indeed possible to accurately recover $\mathbf{x}$ from the ill-posed linear system [3, 4] when the number of measurements M is large enough, and the noise level is modest. However, we might need $M > N$ when the signal is dense or the noise is substantial. Hence, we do not constrain ourselves to the case of $M \ll N$.

Part of the work has appeared in Han et al. [1] and Zhu et al. [2].

The work was supported by the National Science Foundation under the Grants CCF-1217749 and ECCS-1611112.

Junan Zhu and Dror Baron are with the Department of Electrical and Computer Engineering, NC State University, Raleigh, NC 27695. E-mail: {jzhu9, barondror}@ncsu.edu. Ahmad Beirami is with the Research Laboratory of Electronics, Massachusetts Institute of Technology. E-mail: beirami@mit.edu.

Approximate message passing (AMP) [5–8] is an iterative framework that solves linear inverse problems by successively decoupling [9–11] the problem in (1) into scalar denoising problems with additive white Gaussian noise (AWGN). AMP has received considerable attention, because of its fast convergence and the state evolution (SE) formalism [5, 7, 8], which offers a precise characterization of the AWGN denoising problem in each iteration. In the Bayesian setting, AMP often achieves the minimum mean squared error (MMSE) [12–15] in the limit of large linear systems ($N \to \infty$, $\frac{M}{N} \to \kappa$).

In real-world applications, a multi-processor (MP) version of the linear system could be of interest, due to either storage limitations in each individual processor node, or the need for fast computation. This paper considers multi-processor linear systems (MP-LS) [1, 16–20], in which there are P *processor nodes* and a *fusion center*. Each processor node stores $\frac{M}{P}$ rows of the matrix $\mathbf{A}$, and acquires the corresponding linear measurements of the underlying signal $\mathbf{x}$. Without loss of generality, we model the measurement system in processor node $p \in \{1, ..., P\}$ as

$$y_i = \mathbf{A}_i\mathbf{x} + z_i,\ i \in \left\{\frac{M(p-1)}{P} + 1, ..., \frac{Mp}{P}\right\}, \tag{2}$$

where $\mathbf{A}_i$ is the i-th row of $\mathbf{A}$, and y_i and z_i are the i-th entries of $\mathbf{y}$ and $\mathbf{z}$, respectively. Once every y_i is collected, we run distributed algorithms among the fusion center and P processor nodes to estimate the signal $\mathbf{x}$. MP versions of AMP (MP-AMP) for MP-LS have been studied in the literature [1, 18]. Usually, MP platforms are designed for distributed settings such as sensor networks [21, 22] or large-scale "big data" computing systems [23], where the computational and communication burdens can differ among different settings. We reduce the communication costs of MP platforms by applying lossy compression [24–26] to the communication portion of MP-AMP. Our key idea in this work is to minimize the total communication and computation costs by varying the lossy compression schemes in different iterations of MP-AMP.

B. Contribution and organization

Rate-distortion (RD) theory suggests that we can transmit data with greatly reduced coding rates, if we allow some distortion at the output. However, the MP-AMP problem does not directly fall into the RD framework, because the quantization error in the current iteration feeds into estimation errors in future iterations. We quantify the interaction between these two forms of error by studying the excess mean squared error (EMSE) of MP-AMP above the MMSE (EMSE=MSE-MMSE, where MSE denotes the mean squared error). Our

first contribution (Section III) is to use dynamic programming (DP, cf. Bertsekas [27]) to find a sequence of coding rates that yields a desired EMSE while achieving the smallest combined cost of communication and computation; our DP-based scheme is proved to yield optimal coding rates.

Our second contribution (Section IV) is to pose the task of finding the optimal coding rate at each iteration in the low EMSE limit as a convex optimization problem. We prove that the optimal coding rate grows approximately linearly in the low EMSE limit. At the same time, we also provide the theoretic asymptotic growth rate of the optimal coding rates in the limit of low EMSE. This provides practitioners with a heuristic to find a near-optimal coding rate sequence without solving the optimization problem. The linearity of the optimal coding rate sequence (defined in Section III) is also illustrated numerically. With the rate being approximately linear, we obtain that the combined cost of computation and communication scales as $O(\log^2(1/\text{EMSE}))$.

In Section V, we further consider a rich design space that includes various costs, such as the number of iterations T, aggregate coding rate R_{agg}, which is the sum of the coding rates in all iterations and is formally defined in (14), and the MSE achieved by the estimation algorithm. In such a rich design space, reducing any cost is likely to incur an increase in other costs, and it is impossible to simultaneously minimize all the costs. Han et al. [18] reduce the communication costs, and Ma et al. [28] develop an algorithm with reduced computation; both works [18, 28] achieve a reasonable MSE. However, the optimal trade-offs in this rich design space have not been studied. Our third contribution is to pose the problem of finding the best trade-offs among the individual costs T, R_{agg}, and MSE as a multi-objective optimization problem (MOP), and study the properties of Pareto optimal tuples [29] of this MOP. These properties are verified numerically using the DP-based scheme developed in this paper.

Finally, we emphasize that although this paper is presented for the specific framework of MP-AMP, similar methods could be applied to other iterative distributed algorithms, such as consensus averaging [30, 31], to obtain the optimal coding rate as well as optimal trade-offs between communication and computation costs.

The rest of the paper is organized as follows. Section II provides background content. Section III formulates a DP scheme that finds an optimal coding rate. Section IV proves that any optimal coding rate in the low EMSE limit grows approximately linearly as iterations proceed. Section V studies the optimal trade-offs among the computation cost, communication cost, and the MSE of the estimate. Section VI uses some real-world examples to showcase the different trade-offs between communication and computation costs, and Section VII concludes the paper.

II. Background

A. Centralized linear system using AMP

In our linear system (1), we consider an independent and identically distributed (i.i.d.) Gaussian measurement matrix $\mathbf{A}$, i.e., $A_{i,j} \sim \mathcal{N}(0, \frac{1}{M})$, where $\mathcal{N}(\mu, \sigma^2)$ denotes a Gaussian distribution with mean μ and variance σ^2. The signal entries follow an i.i.d. distribution, $f_X(x)$. The noise entries obey $z_i \sim \mathcal{N}(0, \sigma_Z^2)$, where σ_Z^2 is the noise variance.

Starting from $\mathbf{x}_0 = \mathbf{0}$, the AMP framework [5] proceeds iteratively according to

$$\mathbf{x}_{t+1} = \eta_t(\mathbf{A}^\top \mathbf{r}_t + \mathbf{x}_t), \tag{3}$$

$$\mathbf{r}_t = \mathbf{y} - \mathbf{A}\mathbf{x}_t + \frac{1}{\kappa}\mathbf{r}_{t-1}\langle \eta'_{t-1}(\mathbf{A}^\top \mathbf{r}_{t-1} + \mathbf{x}_{t-1})\rangle, \tag{4}$$

where $\eta_t(\cdot)$ is a denoising function, $\eta'_t(\cdot) = \frac{d\eta_t(\cdot)}{d\{\cdot\}}$ is the derivative of $\eta_t(\cdot)$, and $\langle \mathbf{u} \rangle = \frac{1}{N}\sum_{i=1}^N u_i$ for any vector $\mathbf{u} \in \mathbb{R}^N$. The subscript t represents the iteration index, $\{\cdot\}^\top$ denotes the matrix transpose operation, and $\kappa = \frac{M}{N}$ is the measurement rate. Owing to the decoupling effect [9–11], in each AMP iteration [6–8], the vector $\mathbf{f}_t = \mathbf{A}^\top \mathbf{r}_t + \mathbf{x}_t$ in (3) is statistically equivalent to the input signal $\mathbf{x}$ corrupted by AWGN $\mathbf{w}_t$ generated by a source $W \sim \mathcal{N}(0, \sigma_t^2)$,

$$\mathbf{f}_t = \mathbf{x} + \mathbf{w}_t. \tag{5}$$

We call (5) the *equivalent scalar channel*. In large systems ($N \to \infty$, $\frac{M}{N} \to \kappa$),[1] a useful property of AMP [6–8] is that the noise variance σ_t^2 evolves following state evolution (SE):

$$\sigma_{t+1}^2 = \sigma_Z^2 + \frac{1}{\kappa}\text{MSE}(\eta_t, \sigma_t^2), \tag{6}$$

where $\text{MSE}(\eta_t, \sigma_t^2) = \mathbb{E}_{X,W}\left[(\eta_t(X+W) - X)^2\right]$, $\mathbb{E}_{X,W}(\cdot)$ is expectation with respect to (w.r.t.) X and W, and X is the source that generates $\mathbf{x}$. Note that $\sigma_1^2 = \sigma_Z^2 + \frac{\mathbb{E}[X^2]}{\kappa}$, because of the all-zero initial estimate for $\mathbf{x}$. Formal statements for SE appear in prior work [6–8].

In this paper, we confine ourselves to the Bayesian setting, in which we assume knowledge of the true prior, $f_X(x)$, for the signal $\mathbf{x}$. Therefore, throughout this paper we use conditional expectation, $\eta_t(\cdot) = \mathbb{E}[\mathbf{x}|\mathbf{f}_t]$, as the MMSE-achieving denoiser.[2] The derivative of $\eta_t(\cdot)$, which is continuous, can be easily obtained, and is omitted for brevity. Other denoisers such as soft thresholding [5–7] yield MSEs that are larger than that of the MMSE denoiser, $\eta_t(\cdot) = \mathbb{E}[\mathbf{x}|\mathbf{f}_t]$. When the true prior for $\mathbf{x}$ is unavailable, parameter estimation techniques can be used [33]; Ma et al. [34] study the behavior of AMP when the denoiser uses a mismatched prior.

B. MP-LS using lossy MP-AMP

In the sensing problem formulated in (2), the measurement matrix is stored in a distributed manner in each processor node. Lossy MP-AMP [1] iteratively solves MP-LS using lossily compressed messages:

$$\text{Processor nodes: } \mathbf{r}_t^p = \mathbf{y}^p - \mathbf{A}^p\mathbf{x}_t + \frac{1}{\kappa}\mathbf{r}_{t-1}^p\omega_{t-1}, \tag{7}$$

$$\mathbf{f}_t^p = \frac{1}{P}\mathbf{x}_t + (\mathbf{A}^p)^\top \mathbf{r}_t^p, \tag{8}$$

[1] Note that the results of this paper only hold for large systems.

[2] Tan et al. [32] showed that AMP with MMSE-achieving denoisers can be used as a building block for algorithms that minimize arbitrary user-defined error metrics.

$$\text{Fusion center: } \mathbf{f}_{Q,t} = \sum_{p=1}^{P} Q(\mathbf{f}_t^p),\ \omega_t = \langle d\eta_t(\mathbf{f}_{Q,t})\rangle, \quad (9)$$

$$\mathbf{x}_{t+1} = \eta_t(\mathbf{f}_{Q,t}), \quad (10)$$

where $Q(\cdot)$ denotes quantization, and an MP-AMP iteration refers to the process from (7) to (10). The processor nodes send quantized (lossily compressed) messages, $Q(\mathbf{f}_t^p)$, to the fusion center. The reader might notice that the fusion center also needs to transmit the denoised signal vector $\mathbf{x}_t$ and a scalar ω_{t-1} to the processor nodes. The transmission of ω_{t-1} is negligible, and the fusion center may broadcast $\mathbf{x}_t$ so that naive compression of $\mathbf{x}_t$, such as compression with a fixed quantizer, is sufficient. Hence, we will not discuss possible compression of messages transmitted by the fusion center.

Assume that we quantize $\mathbf{f}_t^p, \forall p$, and use C bits to encode the quantized vector $Q(\mathbf{f}_t^p) \in \mathbb{R}^N$. The *coding rate* is $R = \frac{C}{N}$. We incur an *expected distortion*

$$D_t^p = \mathbb{E}\left[\frac{1}{N}\sum_{i=1}^{N}(Q(f_{t,i}^p) - f_{t,i}^p)^2\right]$$

at iteration t in each processor node,[3] where $Q(f_{t,i}^p)$ and $f_{t,i}^p$ are the i-th entries of the vectors $Q(\mathbf{f}_t^p)$ and $\mathbf{f}_t^p$, respectively, and the expectation is over $\mathbf{f}_t^p$. When the size of the problem grows, i.e., $N \to \infty$, the rate-distortion (RD) function, denoted by $R(D)$, offers the fundamental information theoretic limit on the coding rate R for communicating a long sequence up to distortion D [24–26, 35]. A pivotal conclusion from RD theory is that coding rates can be greatly reduced even if D is small. The function $R(D)$ can be computed in various ways [36–38], and can be achieved by an RD-optimal quantization scheme in the limit of large N. Other quantization schemes may require larger coding rates to achieve the same expected distortion D.

The goal of this paper is to understand the fundamental trade-offs for MP-LS using MP-AMP. Hence, unless otherwise stated, we assume that appropriate vector quantization (VQ) schemes [26, 39, 40], which achieve $R(D)$, are applied within each MP-AMP iteration, although our analysis is readily extended to practical quantizers such as entropy coded scalar quantization (ECSQ) [25, 26]. (Note that the cost of running quantizers in each processor node is not considered, because the cost of processing a bit is usually much smaller than the cost of transmitting it.) Therefore, the signal *at the fusion center* before denoising can be modeled as

$$\mathbf{f}_{Q,t} = \sum_{p=1}^{P} Q(\mathbf{f}_t^p) = \mathbf{x} + \mathbf{w}_t + \mathbf{n}_t, \quad (11)$$

where $\mathbf{w}_t$ is the equivalent scalar channel noise (5) and $\mathbf{n}_t$ is the overall quantization error whose entries follow $\mathcal{N}(0, PD_t)$. Because the quantization error, $\mathbf{n}_t$, is a sum of quantization errors in the P processor nodes, $\mathbf{n}_t$ resembles Gaussian noise due to the central limit theorem. Han et al. suggest that SE for lossy MP-AMP [1] (called lossy SE) follows

$$\sigma_{t+1}^2 = \sigma_Z^2 + \frac{1}{\kappa}\text{MSE}(\eta_t, \sigma_t^2 + PD_t), \quad (12)$$

where σ_t^2 can be estimated by $\widehat{\sigma}_t^2 = \frac{1}{M}\|\mathbf{r}_t\|_2^2$ with $\|\cdot\|_p$ denoting the ℓ_p norm [6, 7], and σ_{t+1}^2 is the variance of $\mathbf{w}_{t+1}$.

The rigorous justification of (12) by extending the framework put forth by Bayati and Montanari [7] and Rush and Venkataramanan [8] is left for future work. Instead, we argue that lossy SE (12) asymptotically tracks the evolution of σ_t^2 in lossy MP-AMP in the limit of $\frac{PD_t}{\sigma_t^2} \to 0$. Our argument is comprised of three parts: (*i*) $\mathbf{w}_t$ and $\mathbf{n}_t$ (11) are approximately independent in the limit of $\frac{PD_t}{\sigma_t^2} \to 0$, (*ii*) $\mathbf{w}_t + \mathbf{n}_t$ is approximately independent of $\mathbf{x}$ in the limit of $\frac{PD_t}{\sigma_t^2} \to 0$, and (*iii*) lossy SE (12) holds if (*i*) and (*ii*) hold. The first part ($\mathbf{w}_t$ and $\mathbf{n}_t$ are independent) ensures that we can track the variance of $\mathbf{w}_t + \mathbf{n}_t$ with $\sigma_t^2 + PD_t$. The second part ($\mathbf{w}_t + \mathbf{n}_t$ is independent of $\mathbf{x}$) ensures that lossy MP-AMP follows lossy SE (12) as it falls under the general framework discussed in Bayati and Montanari [7] and Rush and Venkataramanan [8]. Hence, the third part of our argument holds. The first two parts are backed up by extensive numerical evidence in Appendix A, where ECSQ [25, 26] is used; ECSQ approaches $R(D)$ within 0.255 bits in the high rate limit (corresponds to small distortion) [26]. Furthermore, Appendix B provides extensive numerical evidence to show that lossy SE (12) indeed tracks the evolution of the MSE when $\mathbf{w}_t$ and $\mathbf{n}_t$ are independent and $\mathbf{w}_t + \mathbf{n}_t$ and $\mathbf{x}$ are independent.

Although lossy SE (12) requires $\frac{PD_t}{\sigma_t^2} \to 0$, if scalar quantization is used in a practical implementation, then lossy SE approximately holds when $\gamma < \frac{2\sigma_t}{\sqrt{P}}$, where γ is the quantization bin size of the scalar quantizer (details in Appendices A and B). Note that the condition $\gamma < \frac{2\sigma_t}{\sqrt{P}}$ is motivated by Widrow and Kollár [41]. If appropriate VQ schemes [26, 39, 40] are used, then we might need milder requirements than $\frac{PD_t}{\sigma_t^2} \to 0$ in the scalar quantizer case, in order for $\mathbf{w}_t$ and $\mathbf{n}_t$ to be independent and for $\mathbf{w}_t + \mathbf{n}_t$ and $\mathbf{x}$ to be independent.

Denote the coding rate used to transmit $Q(\mathbf{f}_t^p)$ at iteration t by R_t. The sequence $\mathbf{R} = (R_1, ..., R_T)$ is called the *coding rate sequence*, where T is the total number of MP-AMP iterations. Given $\mathbf{R}$, the distortion D_t can be evaluated with $R(D)$, and the scalar channel noise variance σ_t^2 can be evaluated with (12). Hence, the MSE for $\mathbf{R}$ can be predicted. The MSE at the last iteration is called the *final MSE*.

III. Optimal Rates Using Dynamic Programming

In this section, we first define the cost of running MP-AMP. We then use DP to find an optimal coding rate sequence with minimum cost, while achieving a desired EMSE.

Definition 1 (Combined cost): Define the cost of estimating a signal in an MP system as

$$C^b(\mathbf{R}) = b\|\mathbf{R}\|_0 + \|\mathbf{R}\|_1, \quad (13)$$

where $\|\mathbf{R}\|_0 = T$ is the number of iterations to run, and $\|\mathbf{R}\|_1$

[3] Because we assume that $\mathbf{A}$ and $\mathbf{z}$ are both i.i.d., the expected distortions are the same over all P nodes, and can be denoted by D_t for simplicity. Note also that $D_t = \mathbb{E}[(Q(f_{t,i}^p) - f_{t,i}^p)^2]$ due to $\mathbf{x}$ being i.i.d.

is the aggregate coding rate, denoted also by R_{agg},

$$R_{agg} = \|\mathbf{R}\|_1 = \sum_{t=1}^{T} R_t. \tag{14}$$

The parameter b is the cost of computation in one MP-AMP iteration normalized by the cost of transmitting $Q(\mathbf{f}_t^p)$ (9) at a coding rate of 1 bit/entry. Also, the cost at iteration t is

$$C_t^b(R_t) = b \times \mathbb{1}_{R_t \neq 0} + R_t, \tag{15}$$

where the indicator function $\mathbb{1}_{\mathcal{A}}$ is 1 if the condition $\mathcal{A}$ is met, else 0. Hence, $C^b(\mathbf{R}) = \sum_{t=1}^{T} C_t^b(R_t)$.

In some applications, we may want to obtain a sufficiently small EMSE at minimum cost (13), where the physical meaning of the cost varies in different problems (cf. Section VI). Denote the EMSE at iteration t by ϵ_t. Hence, the *final EMSE* at the output of MP-AMP is ϵ_T.

Let us formally state the problem. Our goal is to obtain a coding rate sequence $\mathbf{R}$ for MP-AMP iterations, which is the solution of the following optimization problem:

$$\text{minimize } C^b(\mathbf{R}) \qquad \text{subject to } \epsilon_T \leq \Delta. \tag{16}$$

We now have a definition for the optimal coding rate sequence.

Definition 2 (Optimal coding rate sequence): An optimal coding rate sequence $\mathbf{R}^*$ is a solution of (16).

To compute $\mathbf{R}^*$, we derive a dynamic programming (DP) [27] scheme, and then prove that it is optimal.

Dynamic programming scheme: Suppose that MP-AMP is at iteration t. Define the smallest cost for the $(T-t)$ remaining iterations to achieve the EMSE constraint, $\epsilon_T \leq \Delta$, as $\Phi_{T-t}(\sigma_t^2)$, which is a function of the scalar channel noise variance at iteration t, σ_t^2 (11). Hence, $\Phi_{T-1}(\sigma_1^2)$ is the cost for solving (16), where $\sigma_1^2 = \sigma_Z^2 + \frac{1}{\kappa}\mathbb{E}[X^2]$ is due to the all-zero initialization of the signal estimate.

DP uses a base case and recursion steps to find $\Phi_{T-1}(\sigma_1^2)$. In the base case of DP, $T-t=0$, the cost of running MP-AMP is $C_T^b(R_T) = b \times \mathbb{1}_{R_T \neq 0} + R_T$ (15). If σ_T^2 is not too large, then there exist some values for R_T that satisfy $\epsilon_T \leq \Delta$; for these σ_T^2 and R_T, we have $\Phi_0(\sigma_T^2) = \min_{R_T} C_T^b(R_T)$. If σ_T^2 is too large, even lossless transmission of $\mathbf{f}_T^p$ during the single remaining MP-AMP iteration (12) does not yield an EMSE that satisfies the constraint, $\epsilon_T \leq \Delta$, and we assign $\Phi_0(\sigma_T^2) = \infty$ for such σ_T^2.

Next, in the recursion steps of DP, we iterate back in time by decreasing t (equivalently, increasing $T-t$),

$$\Phi_{T-t}(\sigma_t^2) = \min_{\widehat{R}} \left\{ C_t^b(\widehat{R}) + \Phi_{T-(t+1)}(\sigma_{t+1}^2(\widehat{R})) \right\}, \tag{17}$$

where $\widehat{R}$ is the coding rate used in the current MP-AMP iteration t, the equivalent scalar channel noise variance at the fusion center is σ_t^2 (11), and $\sigma_{t+1}^2(\widehat{R})$, which is obtained from (12), is the variance of the scalar channel noise (11) in the next iteration after transmitting $\mathbf{f}_t^p$ at rate $\widehat{R}$. The terms on the right hand side are the current cost of MP-AMP (15) (including computational and communication costs) and the minimum combined cost in all later iterations, $t+1, ..., T$.

The coding rates $\widehat{R}$ that yield the smallest cost $\Phi_{T-t}(\sigma_t^2)$ for different t and σ_t^2 are stored in a table $\mathcal{R}(t, \sigma_t^2)$. After DP finishes, we obtain the coding rate for the first MP-AMP iteration as $R_1 = \mathcal{R}(1, \sigma_Z^2 + \frac{1}{\kappa}\mathbb{E}[X^2])$. Using R_1, we calculate σ_t^2 from (12) for $t=2$ and find $R_2 = \mathcal{R}(2, \sigma_2^2)$. Iterating from $t=1$ to T, we obtain $\mathbf{R} = (R_1, \cdots, R_T)$.

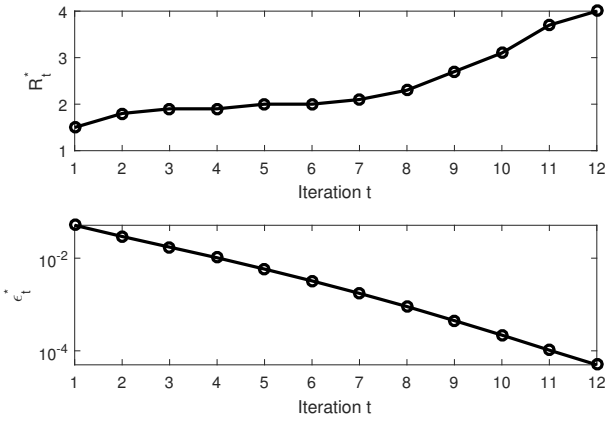

Fig. 1. The optimal coding rate sequence $\mathbf{R}^*$ (top panel) and optimal EMSE ϵ_t^* (bottom) given by DP are shown as functions of t. (Bernoulli Gaussian signal (18) with $\rho = 0.1$, $\kappa = 0.4$, $P = 100$, $\sigma_Z^2 = \frac{1}{400}$, and $b = 2$.)

To be computationally tractable, the proposed DP scheme should operate in discretized search spaces for $\sigma_{\{\cdot\}}^2$ and $R_{\{\cdot\}}$. Details about the resolutions of $\sigma_{\{\cdot\}}^2$ and $R_{\{\cdot\}}$ appear in Appendix C.

In the following, we state that our DP scheme yields the optimal solution. The proof appears in Appendix D.

Lemma 1: The dynamic programming formulation in (17) yields an optimal coding rate sequence $\mathbf{R}^*$, which is a solution of (16) for the discretized search spaces of R_t and σ_t^2, $\forall t$.

Lemma 1 focuses on the optimality of our DP scheme in discretized search spaces for R_t and σ_t^2. It can be shown that we can achieve a desired accuracy level in $\mathbf{R}^*$ by adjusting the resolutions of the discretized search spaces for R_t and σ_t^2. Suppose that the discretized search spaces for $\sigma_{\{\cdot\}}^2$ and $R_{\{\cdot\}}$ have K_1 and K_2 different values, respectively. Then, the computational complexity of our DP scheme is $O(TK_1K_2)$.

Optimal coding rate sequence given by DP: Consider estimating a *Bernoulli Gaussian* signal,

$$X = X_B X_G, \tag{18}$$

where $X_B \sim \text{Ber}(\rho)$ is a Bernoulli random variable, ρ is called the *sparsity rate* of the signal, and $X_G \sim \mathcal{N}(0,1)$; here we use $\rho = 0.1$. Note that the results in this paper apply to priors, $f_X(x)$, other than (18).

We run our DP scheme on a problem with relatively small desired EMSE, $\Delta = 5 \times 10^{-5}$, in the last iteration T. The signal is measured in an MP platform with $P = 100$ processor nodes according to (2). The measurement rate is $\kappa = \frac{M}{N} = 0.4$, and the noise variance is $\sigma_Z^2 = \frac{1}{400}$. The parameter $b = 2$ (13). We use ECSQ [25, 26] as the quantizer in each processor node, and use the corresponding relation between the rate R_t and distortion D_t of ECSQ in our DP scheme. Note that we require the quantization bin size to be smaller than $\frac{2\sigma_t}{\sqrt{P}}$, according to Section II-B. Fig. 1 illustrates the optimal coding rate sequence $\mathbf{R}^*$ and optimal EMSE ϵ_t^*

given by DP as functions of the iteration number t.

It is readily seen that after the first 5–6 iterations the coding rate seems near-linear. The next section proves that any optimal coding rate sequence $\mathbf{R}^*$ is approximately linear in the limit of EMSE$\to 0$. However, our proof involves the large t limit, and does not provide insights for small t. We ran DP for various configurations. Examining all $\mathbf{R}^*$ from our DP results, we notice that the coding rate is monotone non-decreasing, i.e., $R_1^* \leq R_2^* \leq \cdots \leq R_T^*$. This seems intuitive, because in early iterations of (MP-)AMP, the scalar channel noise $\mathbf{w}_t$ is large, which does not require transmitting $\mathbf{f}_t^p$ (cf. (8)) at high fidelity. Hence, a low rate R_t^* suffices. As the iterations proceed, the scalar channel noise $\mathbf{w}_t$ in (11) decreases, and the large quantization error $\mathbf{n}_t$ would be unfavorable for the final MSE. Hence, higher rates are needed in later iterations.

IV. Properties of Optimal Coding Rate Sequences

A. Intuition

We start this section by providing some brief intuitions about why optimal coding rate sequences are approximately linear when the EMSE is small.

Consider a case where we aim to reach a low EMSE. Montanari [6] provided a geometric interpretation of the relation between the MSE performance of AMP at iteration t and the denoiser $\eta_t(\cdot)$ being used.[4] In the limit of small EMSE, the EMSE decreases by a nearly-constant multiplicative factor per AMP iteration, yielding a geometric decay of the EMSE. In MP-AMP, in addition to the equivalent scalar channel noise $\mathbf{w}_t$, we have additive quantization error $\mathbf{n}_t$ (11). In order for the EMSE in an MP-AMP system to decay geometrically, the distortion D_t must decay at least as quickly. To obtain this geometric decay in D_t, recall that in the high rate limit, the distortion-rate function typically takes the form $D(R) \approx C_1 2^{-2R}$ [42] for some positive constant C_1. We propose for R_t to have the form, $R_t \approx C_2 + C_3 t$, where C_2 and C_3 are constants. In the remainder of this section, we first discuss the geometric interpretation of AMP state evolution, followed by our results about the linearity of optimal coding rate sequences. The detailed proofs appear in the appendices.

B. Geometric interpretation of AMP state evolution

Centralized SE: The equivalent scalar channel of AMP is given by (5). We re-write the centralized AMP SE (6) as follows [5, 7, 8],

$$\underbrace{\sigma_{t+1}^2 - \sigma_Z^2}_{g_I(\sigma_{t+1}^2)} = \underbrace{\frac{N}{M}\mathrm{MSE}_{\eta_t}(\sigma_t^2)}_{g_S(\sigma_t^2)}, \tag{19}$$

where $\mathrm{MSE}_{\eta_t}(\sigma_t^2)$ denotes the MSE after denoising $\mathbf{f}_t$ (5) using $\eta_t(\cdot)$. The functions $g_I(\cdot)$ and $g_S(\cdot)$ are illustrated in Fig. 2(a) with solid curves; the meanings of I and S will become clear below. We see that $g_I(\sigma_t^2)$ is an affine function with unit slope, whereas $g_S(\sigma_t^2)$ is generally a nonlinear function of σ_t^2 (see Fig. 2(a)). The lines with arrows illustrate the state evolution (SE). Details appear below.

[4]We will also provide such an interpretation in Section IV-B.

In Fig. 2(a), we present a geometric interpretation of SE. The horizontal axis is the scalar channel noise variance σ^2 and the vertical axis represents the scaled MSE, $u = \frac{N}{M}\mathrm{MSE}$. Let $S_t = (\sigma_t^2, u_t)$ be the *state* point that is reached by SE in iteration t. We follow the SE trajectory $S_t \to I_t \to S_{t+1} \to \cdots$ in Fig. 2(a), where $I_t = (\sigma_{t+1}^2, u_t)$ represents the *intermediate* point in the transition between states S_t and S_{t+1} corresponding to iterations t and $t+1$, respectively. Observe that the points S_t and I_t have the same ordinate (u_t), while S_{t+1} and I_t have the same abscissa (σ_{t+1}^2), which are related as $\sigma_{t+1}^2 = g_I^{-1}(u_t)$ and $u_{t+1} = g_S(\sigma_{t+1}^2)$. As t grows, σ_t^2 converges to σ_∞^2, which is the abscissa of the point S_∞. The ordinate of point S_∞ is $u_\infty = \frac{N}{M}\mathrm{MSE}_\infty$, where $\mathrm{MSE}_\infty = \mathrm{MMSE}$. If we stop the algorithm at iteration T, or equivalently at point $S_T = (\sigma_T^2, u_T)$, the corresponding MSE, MSE_T, has an EMSE of $\epsilon_T = \mathrm{MSE}_T - \mathrm{MMSE}$.

In Fig. 2(b), we zoom into the neighborhood of point S_∞. To make the presentation more concise, we vertically offset $g_I(\cdot)$ and $g_S(\cdot)$ by $\frac{N}{M}$MMSE and horizontally offset them by σ_∞^2; we call the resulting functions $\widetilde{g}_I(\cdot)$ and $\widetilde{g}_S(\cdot)$, respectively. Hence, the vertical axis in Fig. 2(b) represents the scaled EMSE, $\widetilde{u} = \frac{N}{M}\mathrm{EMSE} = \frac{N}{M}\epsilon$, and we have $\widetilde{g}_I(\widetilde{\sigma}_t^2) = g_I(\widetilde{\sigma}_t^2 + \sigma_\infty^2) - \frac{N}{M}\mathrm{MMSE}$ and $\widetilde{g}_S(\widetilde{\sigma}_t^2) = g_S(\widetilde{\sigma}_t^2 + \sigma_\infty^2) - \frac{N}{M}\mathrm{MMSE}$. Observe that $\widetilde{g}_I(0) = \widetilde{g}_S(0) = 0$. Additionally, the slope of $\widetilde{g}_I(\widetilde{\sigma}_t^2)$ is $\widetilde{g}_I'(\widetilde{\sigma}_t^2) = 1$, where $\widetilde{g}_I'(\cdot)$ is the first-order derivative of $\widetilde{g}_I(\cdot)$ w.r.t. $\widetilde{\sigma}_t^2$ (Fig. 2(b)). Because the MSE function for the MMSE-achieving denoiser is continuous and differentiable twice [43], we can invoke Taylor's theorem to express

$$\widetilde{g}_S(\widetilde{\sigma}_t^2) = \widetilde{g}_S'(0)\widetilde{\sigma}_t^2 + \frac{1}{2}\widetilde{g}_S''(\zeta_t)\widetilde{\sigma}_t^4, \tag{20}$$

where $\zeta_t \in (0, \widetilde{\sigma}_t^2)$, and $\widetilde{g}_S'(\widetilde{\sigma}_t^2)$ and $\widetilde{g}_S''(\widetilde{\sigma}_t^2)$ are the first- and second-order derivatives of $\widetilde{g}_S(\cdot)$ w.r.t. $\widetilde{\sigma}_t^2$, respectively. Due to continuity and differentiability of the denoising function, $\widetilde{g}_S(\cdot)$ is invertible in a neighborhood around 0, and its inverse is denoted by $\widetilde{g}_S^{-1}(\cdot)$. Invoking Taylor's theorem,

$$\widetilde{g}_S^{-1}(\widetilde{u}_t) = (\widetilde{g}_S^{-1})'(0)\widetilde{u}_t + \frac{1}{2}(\widetilde{g}_S^{-1})''(\zeta_t)\widetilde{u}_t^2, \tag{21}$$

where $\zeta_t \in (0, \widetilde{u}_t)$, and $(\widetilde{g}_S^{-1})'(\widetilde{u}_t)$ and $(\widetilde{g}_S^{-1})''(\widetilde{u}_t)$ are the first- and second-order derivatives of $\widetilde{g}_S^{-1}(\cdot)$ w.r.t. $\widetilde{u}_t$, respectively. When $t \to \infty$, $\widetilde{\sigma}_t^2 \to 0$ and $\widetilde{u}_t \to 0$, and the higher-order terms become $\frac{1}{2}\widetilde{g}_S''(\xi_t)\widetilde{\sigma}_t^4 = O(\widetilde{\sigma}_t^4)$ and $\frac{1}{2}(\widetilde{g}_S^{-1})''(\zeta_t)\widetilde{u}_t^2 = O(\widetilde{u}_t^2)$. In other words, both $\widetilde{g}_S(\widetilde{\sigma}_t^2)$ and $\widetilde{g}_S^{-1}(\widetilde{u}_t)$ become approximately linear functions, as shown in Fig. 2(b). We further denote the slope of $\widetilde{g}_S(0)$ by θ, i.e.,

$$\theta = \widetilde{g}_S'(0) = \frac{1}{(\widetilde{g}_S^{-1})'(0)}. \tag{22}$$

To calculate the slope θ, we first calculate the scalar channel noise variance for point S_∞, σ_∞^2, by using replica analysis [14, 15],[5] and obtain $\theta = g_S'(\sigma_\infty^2) = \widetilde{g}_S'(0)$. Moreover, the slope of $\widetilde{g}_S(0)$ satisfies $\theta = \widetilde{g}_S'(0) \in (0, 1)$; otherwise, the curves $\widetilde{g}_I(\cdot)$ and $\widetilde{g}_S(\cdot)$ would not intersect at point S_∞.

[5]The outcome of replica analysis [14, 15] is close to simulating SE (19) with a large number of iterations.

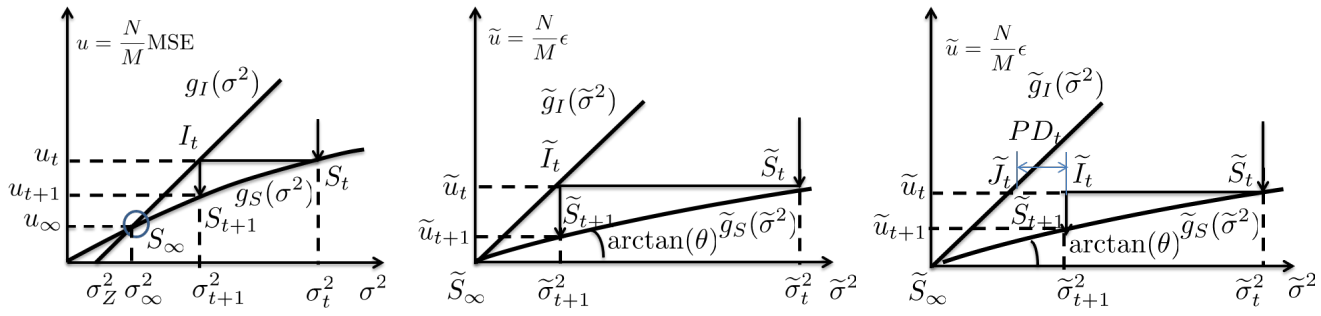

Fig. 2. Geometric interpretation of SE. In all panels, the thick solid curves correspond to $g_I(\cdot)$ and $g_S(\cdot)$, and their offset versions $\widetilde{g}_I(\cdot)$ and $\widetilde{g}_S(\cdot)$. The solid lines with arrows correspond to the SE of AMP. Dashed lines without arrows are auxiliary lines. Panel (a): Illustration of centralized SE. Panel (b): Zooming in to the small region just above point S_∞. Panel (c): Illustration of lossy SE.

Lossy SE: Considering lossy SE (12), we have

$$\underbrace{\sigma_{t+1}^2 - \sigma_Z^2}_{g_I(\sigma_{t+1}^2)} = \underbrace{\frac{N}{M}\text{MSE}_{\eta_t}(\sigma_t^2 + PD_t)}_{g_S(\sigma_t^2 + PD_t)}, \tag{23}$$

where P is the number of processor nodes in an MP network, and D_t is the expected distortion incurred by each node at iteration t. Note that lossy SE has not been rigorously proved in the literature, although we argued in Section II-B that it tracks the evolution of the equivalent scalar channel noise variance σ_t^2 when $D_t \ll \frac{1}{P}\sigma_t^2$.

We notice the additional term PD_t, which corresponds to the distortion *at the fusion center*. Because the P nodes transmit their signals $\mathbf{f}_t^p$ with distortion D_t, and their messages are independent, the fusion center's signal has distortion PD_t. The lines with arrows in Fig. 2(c) illustrate the lossy SE after vertically offsetting $g_I(\cdot)$ and $g_S(\cdot)$ by $\frac{N}{M}$MMSE and horizontally offsetting $g_I(\cdot)$ and $g_S(\cdot)$ by σ_∞^2. After arriving at point $\widetilde{S}_t$, we move horizontally to $\widetilde{J}_t$, and obtain the ordinate of $\widetilde{I}_t$, $\widetilde{u}_t$, from $\widetilde{g}_S(\widetilde{\sigma}_t^2 + PD_t) = \widetilde{u}_t$. Geometrically, SE is dragged to the right by distance PD_t from point $\widetilde{J}_t$ to $\widetilde{I}_t$, and then SE descends from $\widetilde{I}_t$ to $\widetilde{S}_{t+1}$.

C. Asymptotic linearity of the optimal coding rate sequence

Recall from (20) that $\lim_{t\to\infty}\widetilde{\sigma}_t^2 = 0$. Hence, as t grows, $f_{t,i}$ (5) converges in distribution to $x_i + \mathcal{N}(0, \sigma_\infty^2)$. Therefore, the RD function converges to some fixed function as t grows. For large coding rate R, this function has the form

$$R_t = \frac{1}{2}\log_2\left(\frac{C_1}{D_t}\right)(1 + o_t(1)), \tag{24}$$

for some constant C_1 that does not depend on t [42]. Note that the assumption of $\widetilde{\sigma}_t^2$ being small implicitly requires the coding rate used in the corresponding iteration to be large.

For an optimal coding rate sequence $\mathbf{R}^*$, we call the distortion D_t^*, derived from (24), incurred by the optimal code rate R_t^* at a certain iteration t the *optimal distortion*. Correspondingly, we call the EMSE achieved by MP-AMP with $\mathbf{R}^*$, denoted by ϵ_t^*, the *optimal EMSE* at iteration t. In the following, we state our main results on the optimal coding rate, the optimal distortion, and the optimal EMSE.

Theorem 1 (Linearity of the coding rate sequence): Supposing that lossy SE (23) holds, we have

$$\lim_{t\to\infty}\frac{D_{t+1}^*}{D_t^*} = \theta, \tag{25}$$

where θ is defined in (22). Further,

$$\lim_{t\to\infty}\left(R_{t+1}^* - R_t^*\right) = \frac{1}{2}\log_2\left(\frac{1}{\theta}\right). \tag{26}$$

Theorem 1 is proved in Appendix E.

Remark 1: Define the additive growth rate of an optimal coding rate sequence $\mathbf{R}^*$ at iteration t as $R_{t+1}^* - R_t^*$. Theorem 1 not only shows that any optimal coding rate sequence grows approximately linearly in the low EMSE limit, but also provides a way to calculate its additive growth rate in the low EMSE limit. Hence, if the goal is to achieve a low EMSE, practitioners could simply use a coding rate sequence that has a fixed coding rate in the first few iterations and then increases linearly with additive growth rate $\frac{1}{2}\log_2\left(\frac{1}{\theta}\right)$.

The following theorem provides (*i*) the relation between the optimal distortion D_{t+1}^* and the optimal EMSE ϵ_t^* in the large t limit, and (*ii*) the convergence rate of the optimal EMSE ϵ_t^*.

Theorem 2: Assuming that lossy SE (23) holds, we have

$$\lim_{t\to\infty}\frac{D_t^*}{\epsilon_t^*} = 0. \tag{27}$$

Furthermore, the convergence rate of the optimal EMSE is

$$\lim_{t\to\infty}\frac{\epsilon_{t+1}^*}{\epsilon_t^*} = \theta. \tag{28}$$

Theorem 2 is proved in Appendix F. Note that $\lim_{t\to\infty}\frac{D_t^*}{\epsilon_t^*} = 0$ meets the requirement $\frac{PD_t}{\sigma_t^2} \to 0$ discussed in Section II-B. Extending Theorems 1 and 2, we have the following result.

Corollary 3: Assuming that lossy SE (12) holds, the combined computation and communication cost (13) scales as $O(\log^2(1/\Delta))$, $\forall b > 0$, where Δ is the desired EMSE.

Proof: Given Theorem 2, we obtain that the optimal EMSE, ϵ_t^*, indeed decreases geometrically in the large t limit (as a reminder, we provided such intuition in Section IV-A). Considering (14) and Theorem 1, the total computation and communication cost (13) for running T iterations is $C^b(\mathbf{R}^*) = O(T^2) = O(\log^2(1/\epsilon_T^*)) = O(\log^2(1/\Delta))$. ■

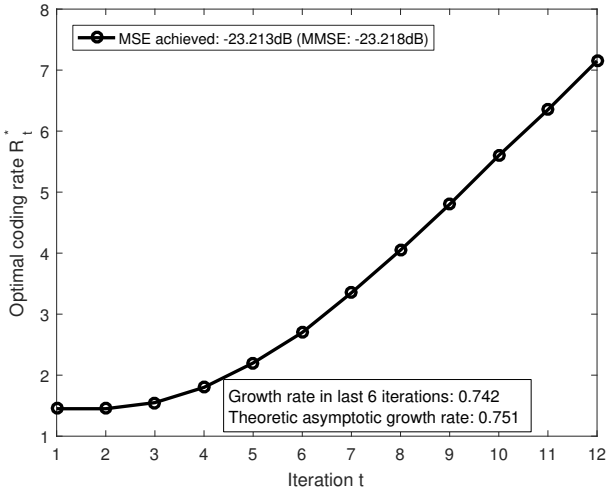
MSE achieved: -23.213dB (MMSE: -23.218dB)
Optimal coding rate R*t
Iteration t
Growth rate in last 6 iterations: 0.742
Theoretic asymptotic growth rate: 0.751

Fig. 3. Comparison of the additive growth rate of the optimal coding rate sequence given by DP at low EMSE and the asymptotic additive growth rate $\frac{1}{2}\log_2\left(\frac{1}{\theta}\right)$. (Bernoulli Gaussian signal (18) with $\rho = 0.2,\ \kappa = 1,\ P = 100, \sigma_Z^2 = 0.01,\ b = 0.782$.)

Remark 2: The key to the proofs of Theorems 1 and 2 is lossy SE (23). We expect that the linearity of the optimal coding rate sequence could be extended to other iterative distributed algorithms provided that (*i*) they have formulations similar to lossy SE (23) that track their estimation errors and (*ii*) their estimation errors converge geometrically. Moreover, formulations that track the estimation error in such algorithms might require less restrictive constraints than AMP. For example, consensus averaging [30, 31] only requires i.i.d. entries in the vector that each node in the network averages.

D. Comparison of DP results to Theorem 1

We run DP (cf. Section III) to find an optimal coding rate sequence $\mathbf{R}^*$ for the setting of $P = 100$ nodes, a Bernoulli Gaussian signal (18) with sparsity rate $\rho = 0.2$, measurement rate $\kappa = 1$, noise variance $\sigma_Z^2 = 0.01$, and parameter $b = 0.782$. The goal is to achieve a desired EMSE of 0.005 dB, i.e., $10\log_{10}\left(1 + \frac{\Delta}{\text{MMSE}}\right) = 0.005$. We use ECSQ [25, 26] as the quantizer in each processor node and use the corresponding relation between the rate R_t and distortion D_t of ECSQ in the DP scheme. Note that we require the quantization bin size γ to be smaller than $\frac{2\sigma_t}{\sqrt{P}}$, according to Section II-B. We know that ECSQ achieves a coding rate within an additive constant of the RD function $R(D)$ [26]. Therefore, the additive growth rate of the optimal coding rate sequence obtained for ECSQ will be the same as the additive growth rate if the RD relation is modeled by $R(D)$ [24–26, 35].

The resulting optimal coding rate sequence is plotted in Fig. 3. The additive growth rate of the last six iterations is $\frac{1}{6}(R_{12}^* - R_6^*) = 0.742$, and the asymptotic additive growth rate according to Theorem 1 is $\frac{1}{2}\log_2\left(\frac{1}{\theta}\right) \approx 0.751$. Note that we use $\Delta R_t = 0.05$ in the discretized search space for R_t. Hence, the discrepancy of 0.009 between the additive growth rate from the simulation and the asymptotic additive growth rate is within our numerical precision. In conclusion, our numerical result matches the theoretical prediction of Theorem 1.

V. Achievable Performance Region

Following the discussion of Section II, we can see that the lossy compression of $\mathbf{f}_t^p, \forall p \in \{1, ..., P\}$, can reduce communication costs. On the other hand, the greater the savings in the coding rate sequence $\mathbf{R}$, the worse the final MSE is expected to be. If a certain level of final MSE is desired despite a small coding rate budget, then more iterations T will be needed. As mentioned above, there is a trade-off between T, R_{agg}, and the final MSE, i.e., MMSE $+ \Delta$, and there is no solution that minimizes them simultaneously. To deal with such trade-offs, which implicitly correspond to sweeping b in (13) in a multi-objective optimization (MOP) problem, it is customary to think about *Pareto optimality* [29].

A. Properties of achievable region

For notational convenience, denote the set of all MSE values achieved by the pair (T, R_{agg}) for some parameter b (13) by $\mathcal{E}(T, R_{agg})$. Within (T, R_{agg}), let the smallest MSE be $\text{MSE}^*(T, R_{agg})$. We now define the achievable set $\mathcal{C}$,

$$\mathcal{C} := \{(T, R_{agg}, \text{MSE}) \in \mathbb{R}_{\geq 0}^3 : \text{MSE} \in \mathcal{E}(T, R_{agg})\},$$

where $\mathbb{R}_{\geq 0}$ is the set of non-negative real numbers. That is, $\mathcal{C}$ contains all tuples (T, R_{agg}, MSE) for which some instantiation of MP-AMP estimates the signal at the desired MSE level using T iterations and aggregate coding rate R_{agg}.

Definition 3: The point $\mathcal{X}_1 \in \mathcal{C}$ is said to dominate another point $\mathcal{X}_2 \in \mathcal{C}$, denoted by $\mathcal{X}_1 \prec \mathcal{X}_2$, if $T_1 \leq T_2$, $R_{agg_1} \leq R_{agg_2}$, and MSE$_1 \leq$ MSE$_2$. A point $\mathcal{X}^* \in \mathcal{C}$ is Pareto optimal if there does not exist $\mathcal{X} \in \mathcal{C}$ satisfying $\mathcal{X} \prec \mathcal{X}^*$. Furthermore, let $\mathcal{P}$ denote the set of all Pareto optimal points,

$$\mathcal{P} := \{\mathcal{X} \in \mathcal{C} : \mathcal{X} \text{ is Pareto optimal}\}. \tag{29}$$

In words, the tuple (T, R_{agg}, MSE) is Pareto optimal if no other tuple $(\widehat{T}, \widehat{R}_{agg}, \widehat{\text{MSE}})$ exists such that $\widehat{T} \leq T$, $\widehat{R}_{agg} \leq R_{agg}$, and $\widehat{\text{MSE}} \leq \text{MSE}$. Thus, the Pareto optimal tuples belong to the boundary of $\mathcal{C}$.

We extend the definition of the number of iterations T to a probabilistic one. To do so, suppose that the number of iterations is drawn from a probability distribution π over $\mathbb{N}$, such that $\sum_{i=1}^{\infty} \pi_i = 1$. Of course, this definition contains a deterministic $T = j$ as a special case with $\pi_j = 1$ and $\pi_i = 0$ for all $i \neq j$. Armed with this definition of Pareto optimality and the probabilistic definition of the number of iterations, we have the following lemma.

Lemma 2: For a fixed noise variance σ_Z^2, measurement rate κ, and P processor nodes in MP-AMP, the achievable set $\mathcal{C}$ is a convex set.

Proof: We need to show that for any $(T^{(1)}, R_{agg}^{(1)}, \text{MSE}^{(1)})$, $(T^{(2)}, R_{agg}^{(2)}, \text{MSE}^{(2)}) \in \mathcal{C}$ and any $0 < \lambda < 1$,

$$\begin{aligned}(\lambda T^{(1)} + (1-\lambda)T^{(2)}, \lambda R_{agg}^{(1)} + (1-\lambda)R_{agg}^{(2)},\\ \lambda \text{MSE}^{(1)} + (1-\lambda)\text{MSE}^{(2)}) \in \mathcal{C}.\end{aligned} \tag{30}$$

This result is shown using time-sharing arguments (see Cover and Thomas [25]). Assume that $(T^{(1)}, R_{agg}^{(1)}, \text{MSE}^{(1)})$, $(T^{(2)}, R_{agg}^{(2)}, \text{MSE}^{(2)}) \in \mathcal{C}$ are achieved by probability distributions $\pi^{(1)}$ and $\pi^{(2)}$, respectively. Let us select all parameters of

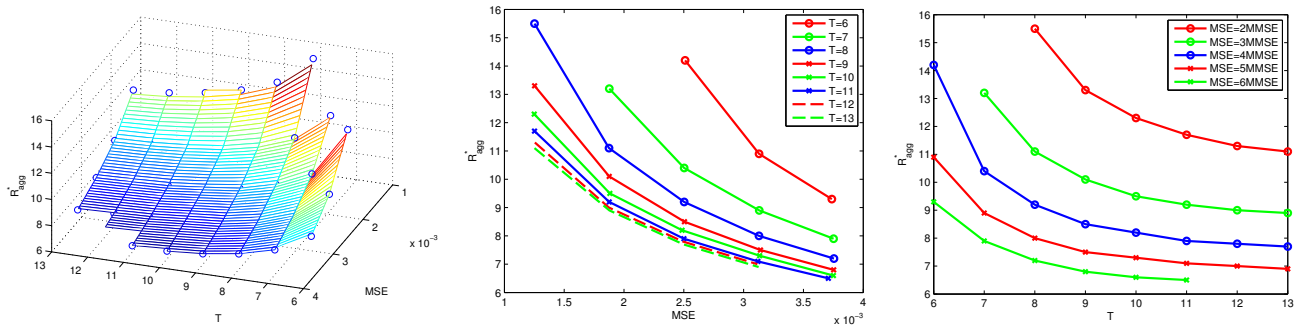

(a) (b) (c)

Fig. 4. Pareto optimal results provided by DP under a variety of parameters b (13): (a) Pareto optimal surface, (b) Pareto optimal aggregate coding rate R^*_{agg} (14) versus the achieved MSE for different optimal MP-AMP iterations T, and (c) Pareto optimal R^*_{agg} (14) versus the number of iterations T for different optimal MSEs. The signal is Bernoulli Gaussian (18) with $\rho = 0.1$. ($\kappa = 0.4$, $P = 100$, and $\sigma_Z^2 = \frac{1}{400}$.)

the first tuple with probability λ and those of the second with probability $(1-\lambda)$. Hence, we have $\pi = \lambda\pi^{(1)} + (1-\lambda)\pi^{(2)}$. Due to the linearity of expectation, $T = \lambda T^{(1)} + (1-\lambda)T^{(2)}$ and $\text{MSE} = \lambda\text{MSE}^{(1)} + (1-\lambda)\text{MSE}^{(2)}$. Again, due to the linearity of expectation, $R_{agg} = \lambda R^{(1)}_{agg} + (1-\lambda)R^{(2)}_{agg}$, implying that (30) is satisfied, and the proof is complete. ■

Definition 4: Let the function $R^*(T, MSE) : \mathbb{R}^2_{\geq 0} \to \mathbb{R}_{\geq 0}$ *be the Pareto optimal rate function, which is implicitly described as* $R^*(T, MSE) = R^*_{agg} \Leftrightarrow (T, R^*_{agg}, MSE) \in \mathcal{P}$. *We further define implicit functions* $T^*(R_{agg}, MSE)$ *and* $MSE^*(T, R_{agg})$ *in a similar way.*

Corollary 4: The functions $R^*(T, \text{MSE})$, $T^*(R_{agg}, \text{MSE})$, and $\text{MSE}^*(T, R_{agg})$ are convex in their arguments.

Note that our proof for the convexity of the set $\mathcal{C}$ might be extended to other iterative distributed learning algorithms that transmit lossily compressed messages.

B. Pareto optimal points via DP

After proving that the achievable set $\mathcal{C}$ is convex, we apply DP in Section III to find the Pareto optimal points, and validate the convexity of the achievable set.

According to Definition 3, the resulting tuple (T, R_{agg}, MSE) computed using DP (Section III) is Pareto optimal on the discretized search spaces. Hence, in this subsection, we run DP to obtain the Pareto optimal points for a certain distributed linear system by sweeping the parameter b (13).

Consider the same setting as in Fig. 1, except that we analyze MP platforms [21–23] for different b (13). Running the DP scheme of Section III, we obtain the optimal coding rate sequence $\mathbf{R}^*$ that yields the lowest combined cost while providing a desired EMSE that is at most $\Delta \in \{1, 2, ..., 5\} \times$ MMSE or equivalently $\text{MSE} \in \{2, 3, ..., 6\} \times$ MMSE. In Fig. 4(a), we draw the Pareto optimal surface obtained by our DP scheme, where the circles are Pareto optimal points. Fig. 4(b) plots the aggregate coding rate R_{agg} as a function of MSE for different optimal numbers of MP-AMP iterations T. Finally, Fig. 4(c) plots the aggregate coding rate R_{agg} as a function of T for different optimal MSEs. We can see that the surface comprised of the Pareto optimal points is indeed convex. Note that when running DP to generate Fig. 4, we used the RD function [24–26, 35] to model the relation between the rate R_t and distortion D_t at each iteration, which could be approached by VQ at sufficiently high rates. We also ignored the constraint on the quantization bin size (Section II-B). Therefore, we only present Fig. 4 for illustration purposes.

When a smaller MSE (or equivalently smaller EMSE) is desired, more iterations T and greater aggregate coding rates R_{agg} (14) are needed. Optimal coding rate sequences increase R_{agg} to reduce T when communication costs are low (examples are commercial cloud computing systems [23], multiprocessor CPUs, and graphic processing units), whereas more iterations allow to reduce the coding rate when communication is costly (for example, in sensor networks [21, 22]). These applications are discussed in Section VI.

Discussion of corner points: We further discuss the corners of the Pareto optimal surface (Fig. 4) below.

1) First, consider the corner points along the MSE coordinate.
   - If $\text{MSE}^* \to$ MMSE (or equivalently $\Delta \to 0$), then MP-AMP needs to run infinite iterations with infinite coding rates. Hence, $R^*_{agg} \to \infty$ and $T^* \to \infty$. The rate of growth of R^*_{agg} can be deduced from Theorem 1.
   - If $\text{MSE}^* \to \rho$ (the variance of the signal (18)), then MP-AMP does not need to run any iterations at all. Instead, MP-AMP outputs an all-zero estimate. Therefore, $\lim_{\text{MSE}^* \to \rho} R^*_{agg} = 0$ and $\lim_{\text{MSE}^* \to \rho} T^* = 0$.
2) Next, we discuss the corner points along the T coordinate.
   - If $T^* \to 0$, then the best MP-AMP can do is to output an all-zero estimate. Hence, $\lim_{T^* \to 0} \text{MSE}^* = \rho$ and $\lim_{T^* \to 0} R^*_{agg} = 0$.
   - The other extreme, $T^* \to \infty$, occurs only when we want to achieve an $\text{MSE}^* \to$ MMSE. Hence, $R_{agg} \to \infty$.
3) We conclude with corner points along the R_{agg} coordinate.
   - If $R^*_{agg} \to 0$, then the best MP-AMP can do is to output an all-zero estimate without running any iterations at all. Hence, $\lim_{R^*_{agg} \to 0} \text{MSE}^* = \rho$ and $\lim_{R^*_{agg} \to 0} T^* = 0$.
   - If $R^*_{agg} \to \infty$, then the optimal scheme will use high rates in all iterations, and MP-AMP resembles centralized AMP. Therefore, the MSE^* as a function of T^* converges to that of centralized AMP SE (6).

VI. Real-world Case Study

To showcase the difference between optimal coding rate sequences in different platforms, this section discusses several MP platforms including sensor networks [21, 22] and large-scale cloud servers [23]. The costs in these platforms are quite different due to the different constraints in these platforms, and we will see how they affect the optimal coding rate sequence $\mathbf{R}^*$. The changes in the optimal $\mathbf{R}^*$ highlight the importance of optimizing for the correct costs.

A. Sensor networks

In sensor networks [21, 22], distributed sensors are typically dispatched to remote locations where they collect data and communicate with the fusion center. However, distributed sensors may have severe power consumption constraints. Therefore, low power chips such as the CC253X from Texas Instruments [44] are commonly used in distributed sensors. Some typical parameters for such low power chips are: central processing unit (CPU) clock frequency 32MHz, data transmission rate 250Kbps, voltage between 2V-3.6V, and transceiver current 25mA [44], where the CPU current resembles the transceiver current. Because these chips are generally designed to be low power, when transmitting and receiving data, the CPU helps the transceiver and cannot carry out computing tasks. Therefore, the power consumption can be viewed as constant. Hence, in order to minimize the power consumption, we minimize the total runtime when estimating a signal from MP-LS measurements (2) collected by the distributed sensors.

The runtime in each MP-AMP iteration (7)-(10) consists of (*i*) time for computing (7) and (8), (*ii*) time for encoding $\mathbf{f}_t^p$ (8), and (*iii*) data transmission time for $Q(\mathbf{f}_t^p)$ (9). As discussed in Section II-B, the fusion center may broadcast $\mathbf{x}_t$ (10), and simple compression schemes can reduce the coding rate. Therefore, we consider the data reception time in the P processor nodes to be constant. The overall computational complexity for (7) and (8) is $O(\frac{MN}{P})$. Suppose further that (*i*) each processor node needs to carry out two matrix-vector products in each iteration, (*ii*) the overhead of moving data in memory is assumed to be 10 times greater than the actual computation, and (*iii*) the clock frequency is 32MHz. Hence, we assume that the actual time needed for computing (7) and (8) is $C_4 = \frac{20MN}{32\times 10^6 P}$ sec. Transmitting $Q(\mathbf{f}_t^p)$ of length N at coding rate R requires $\frac{RN}{250\times 10^3}$ sec, where the denominator is the data transmission rate of the transceiver. Assuming that the overhead in communication is approximately the same as the communication load caused by the actual messages, we obtain that the time requested for transmitting $Q(\mathbf{f}_t^p)$ at coding rate R is $C_5 R$ sec, where $C_5 = \frac{2N}{250\times 10^3}$. Therefore, the total cost can be calculated from (13) with $b = \frac{C_4}{C_5}$ (13).

Because low power chips equipped in distributed sensors have limited memory (around 10KB, although sometimes external flash is allowed) [44], the signal length N and number of measurements M cannot be too large. We consider $N = 1{,}000$ and $M = 400$ spread over $P = 100$ sensors, sparsity rate $\rho = 0.1$, and $\sigma_Z^2 = \frac{1}{400}$. We set the desired MSE to be 0.5 dB above the MMSE, i.e., $10\log_{10}\left(1 + \frac{\Delta}{\text{MMSE}}\right) = 0.5$, and run DP as in Section III.[6] The coding rate sequence provided by DP is $\mathbf{R}^* = (0.1,\ 0.1,\ 0.6,\ 0.8,\ 1.0,\ 1.0,\ 1.1,\ 1.1,\ 1.2,\ 1.4,\ 1.6,\ 1.9,\ 2.3,\ 2.7,\ 3.1)$. In total we have $T = 15$ MP-AMP iterations with $R_{agg} = 20.0$ bits aggregate coding rate (14). The final MSE (MMSE + Δ) is 7.047×10^{-4}, which is 0.5 dB from the MMSE (6.281×10^{-4}) [12–15].

B. Large-scale cloud server

Having discussed sensor networks [21, 22], we now discuss an application of DP (cf. Section III) to large-scale cloud servers. Consider the dollar cost for users of Amazon EC2 [23], a commercial cloud computing service. A typical cost for CPU time is \$0.03/hour, and the data transmission cost is \$0.03/GB. Assuming that the CPU clock frequency is 2.0GHz and considering various overheads, we need a runtime of $\frac{20MN}{2\times 10^9 P}$ sec and the computation cost is $C_4 = \$\frac{20MN}{2\times 10^9 P} \times \frac{0.03}{3600}$ per MP-AMP iteration. Similar to Section VI-A, the communication cost for coding rate R is $C_5 R = \$2RN\frac{0.03}{8\times 10^9}$. Note that the multiplicative factors of 20 in C_4 and 2 in C_5 are due to the same considerations as in Section VI-A, and the 8×10^9 in C_5 is the number of bits per GB. Therefore, the total cost with T MP-AMP iterations can still be modeled as in (13), where $b = \frac{C_4}{C_5}$.

We consider a problem with the same signal and channel model as the setting of Section VI-A, while the size of the problem grows to $N = 50{,}000$ and $M = 20{,}000$ spread over $P = 100$ computing nodes. Running DP, we obtain the coding rate sequence $\mathbf{R}^* = (1.3,\ 1.6,\ 1.8,\ 1.8,\ 1.8,\ 1.9,\ 2.1,\ 2.3,\ 2.6,\ 3.1,\ 3.7)$ for a total of $T = 11$ MP-AMP iterations with $R_{agg} = 24.0$ bits aggregate coding rate. The final MSE is 7.031×10^{-4}, which is 0.49 dB above the MMSE. Note that this final MSE is 0.01 dB better than our goal of 0.5 dB above the MMSE due to the discretized search spaces used in DP.

Settings with even cheaper communication costs: Compared to large-scale cloud servers, the relative cost of communication is even cheaper in multi-processor CPU and graphics processing unit (GPU) systems. We reduce b by a factor of 100 compared to the large-scale cloud server case above. We rerun DP, and obtain the coding rate sequence $\mathbf{R}^* = (2.3,\ 2.5,\ 2.6,\ 2.7,\ 2.7,\ 2.8,\ 3.0,\ 3.4,\ 3.7,\ 4.5)$ for $T = 10$ and $R_{agg} = 30.2$ bits. Note that 10 iterations are needed for centralized AMP to converge in this setting. With the low-cost communication of this setting, DP yields a coding rate sequence $\mathbf{R}^*$ within 0.5 dB of the MMSE with the same number of iterations as centralized AMP, while using an average coding rate of only 3.02 bits per iteration.

Remark 3: Let us review the cost tuples (T, R_{agg}, MSE) for our three cases. For sensor networks, $(T, R_{agg}, \text{MSE})_{\text{sensornet}} = (15, 20, 7.047 \times 10^{-4})$; for cloud servers, $(T, R_{agg}, \text{MSE})_{\text{cloud}} = (11, 24, 7.031 \times 10^{-4})$; and for GPUs, $(T, R_{agg}, \text{MSE})_{\text{GPU}} = (10, 30.2, 7.047 \times 10^{-4})$. These cost tuples are different points in the Pareto optimal set $\mathcal{P}$ (29). We can see for sensor networks that the optimal coding rate

[6] Throughout Section VI, we use the RD function [24–26, 35] to model the relation between rate R_t and distortion D_t at each iteration. We also ignore the constraint on the quantizer (Section II-B). Therefore, the optimal coding rate sequences in Section VI are only for illustration purposes.

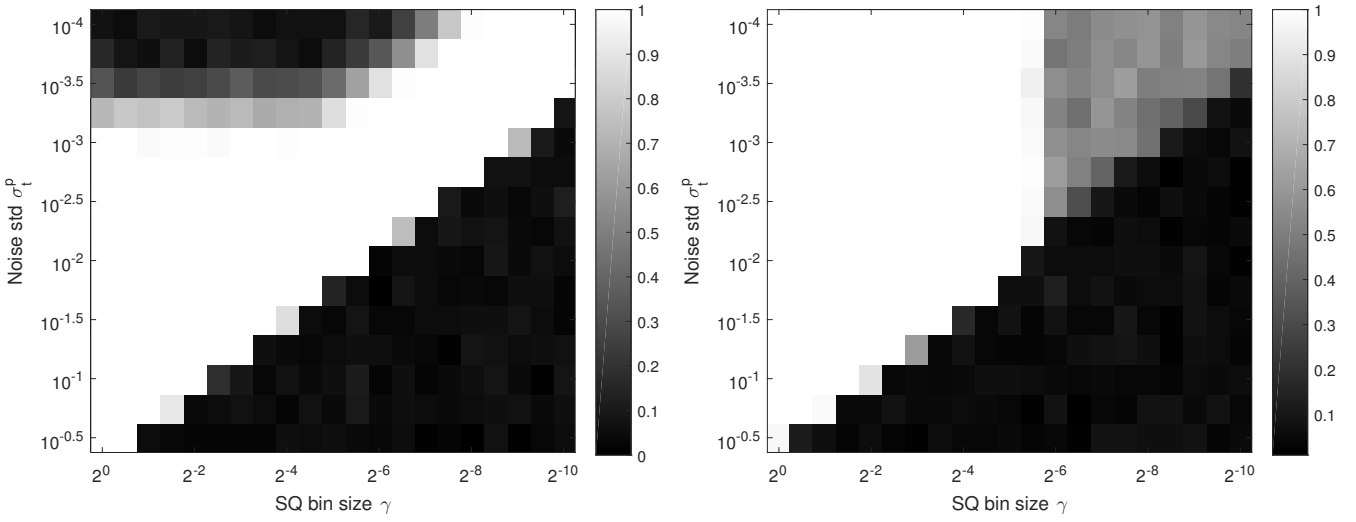

Fig. 5. PCC test results. The shades of gray show the fraction of 100 tests where we reject the null hypothesis (random variables being tested are uncorrelated) with 5% confidence. The horizontal and vertical axes represent the quantization bin size γ of the SQ and the scalar channel noise standard deviation (std) σ_t^p in each processor node, respectively. Panel (a): Test the correlation between $\mathbf{w}_t$ and $\mathbf{n}_t$. Panel (b): Test the correlation between $\mathbf{w}_t + \mathbf{n}_t$ and $\mathbf{x}$.

sequence reduces R_{agg} while adding iterations, because sensor networks have relatively expensive communications. The optimal coding rate sequences use higher rates in cloud servers and GPUs, because their communication costs are relatively lower. Indeed, different trade-offs between computation and communication lead to different aggregate coding rates R_{agg} and numbers of MP-AMP iterations T. Moreover, the optimal coding rate sequences for sensor networks, cloud servers, and GPUs use average coding rates of 1.33, 2.18, and 3.02 bits/entry/iteration, respectively. Compared to 32 bits/entry/iteration single-precision floating point communication schemes, optimal coding rate sequences reduce the communication costs significantly.

VII. Conclusion

This paper used lossy compression in multi-processor (MP) approximate message passing (AMP) for solving MP linear inverse problems. Dynamic programming (DP) was used to obtain the optimal coding rate sequence for MP-AMP that incurs the lowest combined cost of communication and computation while achieving a desired mean squared error (MSE). We posed the problem of finding the optimal coding rate sequence in the low excess MSE (EMSE=MSE-MMSE, where MMSE refers to the minimum MSE) limit as a convex optimization problem and proved that optimal coding rate sequences are approximately linear when the EMSE is small. Additionally, we obtained that the combined cost of computation and communication scales with $O(\log^2(1/\text{EMSE}))$. Furthermore, realizing that there is a trade-off among the communication cost, computation cost, and MSE, we formulated a multi-objective optimization problem (MOP) for these costs and studied the Pareto optimal points that exploit this trade-off. We proved that the achievable region of the MOP is convex.

We further emphasize that there is little work in the prior art discussing the optimization of communication schemes in iterative distributed algorithms. Although we focused on the MP-AMP algorithm, our conclusions such as the linearity of the optimal coding rate sequence and the convexity of the achievable set of communication/computation trade-offs could be extended to other iterative distributed algorithms including consensus averaging [30, 31].

Acknowledgments

The authors thank Puxiao Han and Ruixin Niu for numerous discussions about MP settings of linear systems and AMP. Mohsen Sardari and Mihail Sichitiu offered useful information about the practical distributed computing platforms of Section VI, and Cynthia Rush provided useful insights about the proof of AMP SE. The authors also thank Yanting Ma and Ryan Pilgrim for helpful comments about the manuscript.

Appendix

A. Impact of the quantization error

This appendix provides numerical evidence that (*i*) the quantization error $\mathbf{n}_t$ is independent of the scalar channel noise $\mathbf{w}_t$ (11) in the fusion center and (*ii*) $\mathbf{w}_t + \mathbf{n}_t$ is independent of the signal $\mathbf{x}$. In the following, we simulate the AMP equivalent scalar channel in each processor node and in the fusion center. In the interest of simple implementation, we use scalar quantization (SQ) to quantize $\mathbf{f}_t^p$ (8) (in each processor node) and hypothesis testing to evaluate (*i*) whether $\mathbf{w}_t$ and $\mathbf{n}_t$ (in the fusion center) are independent and (*ii*) whether $\mathbf{w}_t + \mathbf{n}_t$ and $\mathbf{x}$ are independent. Both parts are necessary for lossy SE (12) to hold: part (*i*) ensures that we can predict the variance of $\mathbf{w}_t + \mathbf{n}_t$ by $\sigma_t^2 + PD_t$ and part (*ii*) ensures that lossy MP-AMP falls within the general framework of Bayati and Montanari [7] and Rush and Venkataramanan [8], so that lossy SE (12) holds. Details about our simulation appear below.

Considering (5) and (8), we obtain that the AMP equivalent scalar channel in each processor node can be expressed as

$$\mathbf{f}_t^p = \frac{1}{P}\mathbf{x} + \mathbf{w}_t^p, \tag{31}$$

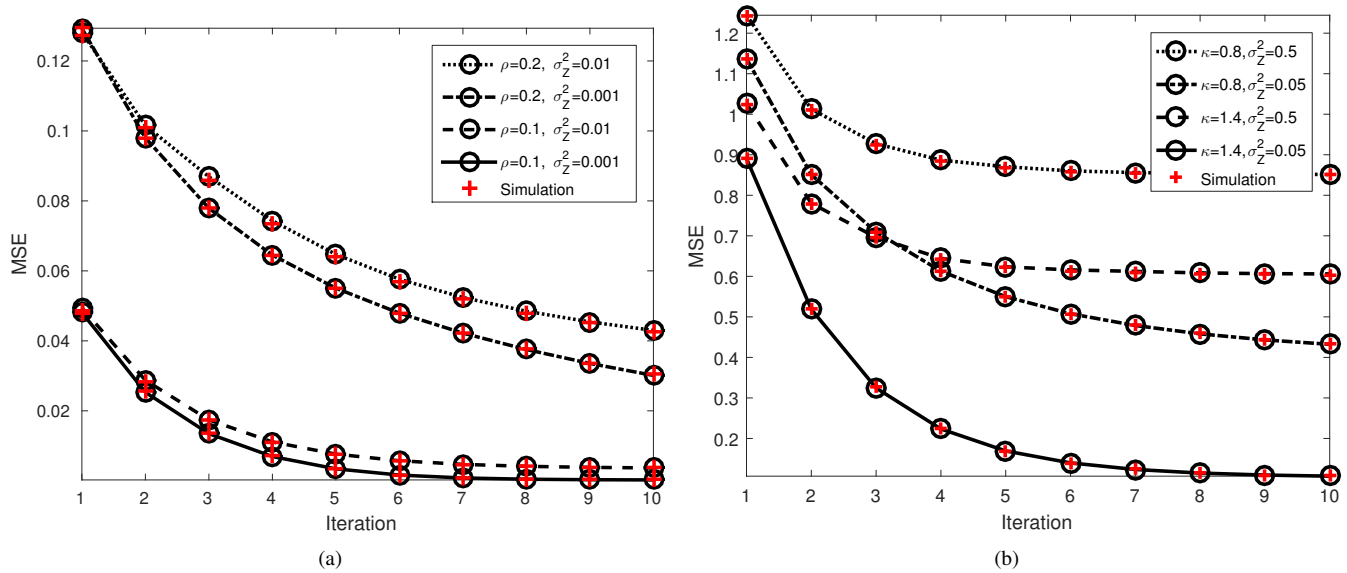

(a) (b)

Fig. 6. Comparison of the MSE predicted by lossy SE (12) and the MSE of MP-AMP simulations for various settings. The round markers represent MSEs predicted by lossy SE, and the (red) crosses represent simulated MSEs. Panel (a): Bernoulli Gaussian signal. Panel (b): Mixture Gaussian signal.

where $\sum_{p=1}^{P} \mathbf{w}_t^p = \mathbf{w}_t$ (5), and the variances of $\mathbf{w}_t^p$ and $\mathbf{w}_t$ can be expressed as $(\sigma_t^p)^2$ and σ_t^2, respectively (12). Hence, we obtain $\sigma_t^2 = \sum_{p=1}^{P} (\sigma_t^p)^2$. The signal $\mathbf{x}$ follows (18) with $\rho = 0.1$. The entries of $\mathbf{w}_t^p$ are i.i.d. and follow $\mathcal{N}(0, (\sigma_t^p)^2)$. Next, we apply an SQ to $\mathbf{f}_t^p$ (31),

$$Q(\mathbf{f}_t^p) = \frac{1}{P}\mathbf{x} + \mathbf{w}_t^p + \mathbf{n}_t^p, \tag{32}$$

where $Q(\cdot)$ denotes the quantization process, $\mathbf{n}_t^p$ is the quantization error in processor node p, and recall that the variance of $\mathbf{n}_t^p$ is D_t. We simulate the fusion center by calculating

$$\mathbf{f}_t = \sum_{p=1}^{P} Q(\mathbf{f}_t^p) = \mathbf{x} + \mathbf{w}_t + \mathbf{n}_t, \tag{33}$$

where $\mathbf{n}_t = \sum_{p=1}^{P} \mathbf{n}_t^p$. Note that $\mathbf{w}_t$ is Gaussian due to properties of AMP [5–7]. The total quantization error at the fusion center, $\mathbf{n}_t$, is also Gaussian, due to the central limit theorem. Hence, in order to test the independence of $\mathbf{w}_t$ and $\mathbf{n}_t$ (33), we need only test whether $\mathbf{w}_t$ and $\mathbf{n}_t$ are uncorrelated. We also test whether $\mathbf{w}_t + \mathbf{n}_t$ and $\mathbf{x}$ are uncorrelated.

We study the settings $\sigma_t^p \in \{10^{-0.5}, \cdots, 10^{-4}\}$ and $\gamma \in \{2^0, \cdots, 2^{-10}\}$, where γ denotes the SQ bin size. In each setting, we simulate (31)–(33) 100 times and perform 100 Pearson correlation coefficient (PCC) tests [45] for $\mathbf{w}_t$ and $\mathbf{n}_t$, respectively. The null hypothesis of the PCC tests [45] is that $\mathbf{w}_t$ and $\mathbf{n}_t$ are uncorrelated. The null hypothesis is rejected if the resulting p-value is smaller than 0.05.

For each setting, we record the fraction of 100 tests where the null hypothesis is rejected, which is shown by the shades of gray in Fig. 5(a). The horizontal and vertical axes represent the quantization bin size γ and the standard deviation (std) σ_t^p, respectively. Similarly, we test $\mathbf{w}_t + \mathbf{n}_t$ and $\mathbf{x}$; results appear in Fig. 5(b). We can see that when $\gamma \ll \sigma_t^p$ (bottom right corner), (*i*) $\mathbf{w}_t$ and $\mathbf{n}_t$ tend to be independent and (*ii*) $\mathbf{w}_t + \mathbf{n}_t$ and $\mathbf{x}$ tend to be independent.

Now consider Fig. 5(b), which provides PCC test results evaluating possible correlations between $\mathbf{w}_t + \mathbf{n}_t$ and $\mathbf{x}$. There appears to be a phase transition that separates regions where $\mathbf{w}_t + \mathbf{n}_t$ and $\mathbf{x}$ seem independent or dependent. We speculate that this phase transition is related to the pdf of $\frac{1}{P}\mathbf{x} + \mathbf{w}_t^p$. To explain our hypothesis, note that when the noise $\mathbf{w}_t^p$ is low (top part of Fig. 5(b)), the phase transition is less affected by noise, and the role of γ is smaller. By contrast, large noise (bottom) sharpens the phase transition.

In summary, it appears that when $\gamma < 2\sigma_t^p = \frac{2\sigma_t}{\sqrt{P}}$, we can regard (*i*) $\mathbf{w}_t$ and $\mathbf{n}_t$ to be independent and (*ii*) $\mathbf{w}_t + \mathbf{n}_t$ and $\mathbf{x}$ to be independent. The requirement $\gamma < 2\sigma_t^p = \frac{2\sigma_t}{\sqrt{P}}$ is motivated by Widrow and Kollár [41]; we leave the study of this phase transition for future work.

B. Numerical evidence for lossy SE

This appendix provides numerical evidence for lossy SE (23). We simulate two signal types, one is the Bernoulli Gaussian signal (18) and the other is a mixture Gaussian.

Bernoulli Gaussian signals: We generate 50 signals of length 10,000 according to (18). These signals are measured by $M = 5{,}000$ measurements spread over $P = 100$ distributed nodes. We estimate each of these signals by running $T = 10$ MP-AMP iterations. ECSQ is used to quantize $\mathbf{f}_t^p$ (31), and $Q(\mathbf{f}_t^p)$ (32) is encoded at coding rate R_t. We simulate settings with sparsity rate $\rho \in \{0.1, 0.2\}$ and noise variance $\sigma_Z^2 \in \{0.01, 0.001\}$. In each setting, we randomly generate the coding rate sequence $\mathbf{R}$, s.t. the quantization bin size at each iteration satisfies $\gamma < \frac{2\sigma_t}{\sqrt{P}}$ (details in Appendix A).[7] A Bayesian denoiser, $\eta_t(\cdot) = \mathbb{E}[\mathbf{x}|\mathbf{f}_t]$, is used in (10). The resulting MSEs from the MP-AMP simulation averaged over the 50 signals, along with MSEs predicted by lossy SE (23),

[7]Note that the constraint on γ implies that $\mathbf{R}$ is likely monotone non-decreasing.

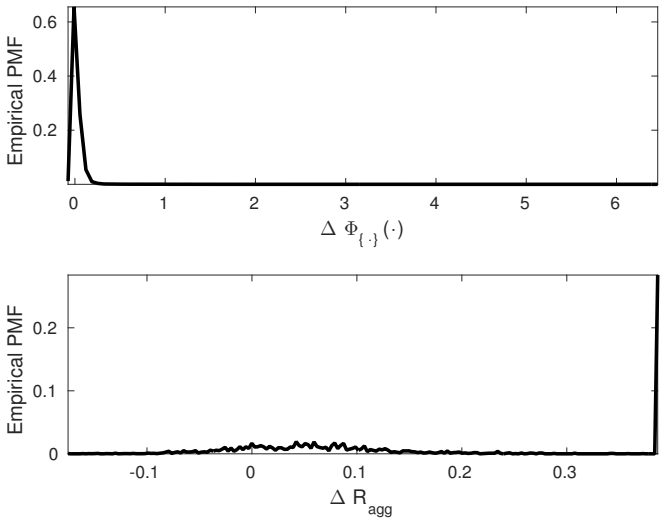

Fig. 7. Justification of the discretized search space used in DP. Top panel: Empirical PMF of the error in the cost function $\Delta\Phi_{\{\cdot\}}(\cdot)$ used to verify the integrity of the linear interpolation in the discretized search space of σ^2. Bottom panel: Empirical PMF of ΔR_{agg}; used to verify the integrity of the choice of $\Delta R = 0.1$.

are plotted in Fig. 6(a). We can see that the simulated MSEs are close to the MSEs predicted by lossy SE.

Mixture Gaussian signals: We independently generate 50 signals of length 10,000 according to $X = \sum_{i\in\{0,1,2\}} \mathbb{1}_{X_B=i} X_{G,i}$ where $X_B \sim \text{cat}(0.5, 0.3, 0.2)$ follows a categorical distribution on alphabet $\{0, 1, 2\}$, $X_{G,0} \sim \mathcal{N}(0, 0.1)$, $X_{G,1} \sim \mathcal{N}(-1.5, 0.8)$, and $X_{G,2} \sim \mathcal{N}(2, 1)$. We simulate settings with $T = 10$, $P = 100$, $\kappa = \frac{M}{N} \in \{0.8, 1.6\}$, and $\sigma_Z^2 \in \{0.5, 0.05\}$. In each setting, we randomly generate the coding rate sequence $\mathbf{R}$, s.t. the quantization bin size at each iteration satisfies $\gamma < \frac{2\sigma_t}{\sqrt{P}}$. The results are plotted in Fig. 6(b). The simulation results match well with the lossy SE predictions.

C. Integrity of discretized search space

When a coding rate $\widehat{R}$ is selected in MP-AMP iteration t, DP calculates the equivalent scalar channel noise variance σ_{t+1}^2 (11) for the next MP-AMP iteration according to (12). The variance σ_{t+1}^2 is unlikely to lie on the discretized search space for σ_t^2, denoted by the *grid* $\mathcal{G}(\sigma^2)$. Therefore, $\Phi_{T-(t+1)}(\sigma_{t+1}^2(\widehat{R}))$ in (17) does not reside in memory. Instead of brute-force calculation of $\Phi_{\{\cdot\}}(\cdot)$, we estimate it by fitting a function to the closest neighbors of σ_{t+1}^2 that lie on the grid $\mathcal{G}(\sigma^2)$ and finding $\Phi_{\{\cdot\}}(\cdot)$ according to the fit function. We evaluate a linear interpolation scheme.

Interpolation in $\mathcal{G}(\sigma^2)$: We run DP over the original coarse grid $\mathcal{G}^c(\sigma^2)$ with resolution $\Delta\sigma^2 = 0.01$ dB, and a $4\times$ finer grid $\mathcal{G}^f(\sigma^2)$ with $\Delta\sigma^2 = 0.0025$ dB. We obtain the cost function with the coarse grid $\Phi_{T-t}^c((\sigma_t^2)_c)$ and the cost function with the fine grid $\Phi_{T-t}^f((\sigma_t^2)_f)$, $\forall t \in \{1, ..., T\}, (\sigma_t^2)_c \in \mathcal{G}^c(\sigma^2), (\sigma_t^2)_f \in \mathcal{G}^f(\sigma^2)$. Next, we interpolate $\Phi_{T-t}^c((\sigma_t^2)_c)$ over the fine grid $\mathcal{G}^f(\sigma^2)$ and obtain the interpolated $\Phi_{T-t}^i((\sigma_t^2)_c)$. In order to compare $\Phi_{T-t}^i((\sigma_t^2)_c)$ with $\Phi_{T-t}^f((\sigma_t^2)_c)$ in a comprehensive way, we consider the settings given by the Cartesian product of the following variables: (*i*) the number of distributed nodes $P \in \{50, 100\}$, (*ii*) sparsity rate $\rho \in \{0.1, 0.2\}$, (*iii*) measurement rate $\kappa = \frac{M}{N} \in \{3\rho, 5\rho\}$, (*iv*) EMSE $\epsilon_T \in \{1, 0.5\}$dB, (*v*) parameter $b \in \{0.5, 2\}$, and (*vi*) noise variance $\sigma_Z^2 \in \{0.01, 0.001\}$. In total, there are 64 different settings. We calculate the error $\Delta\Phi_{T-t}\left((\sigma_t^2)_c\right) = \Phi_{T-t}^i\left((\sigma_t^2)_c\right) - \Phi_{T-t}^f\left((\sigma_t^2)_c\right)$ and plot the empirical probability mass function (PMF) of $\Delta\Phi_{T-t}\left((\sigma_t^2)_c\right)$ over all t, $(\sigma_t^2)_c$, and all 64 settings. The resulting empirical PMF of $\Delta\Phi_{\{\cdot\}}(\cdot)$ is plotted in the top panel of Fig. 7. We see that with 99% probability, the error satisfies $\Delta\Phi_{\{\cdot\}}(\cdot) \leq 0.2$, which corresponds to an inaccuracy of approximately 0.2 in the aggregate coding rate R_{agg}.[8] In the simulation, we used a resolution of $\Delta R = 0.1$. Hence, the inaccuracy of 0.2 in R_{agg} (over roughly 10 iterations) is negligible. Therefore, we use linear interpolation with a coarse grid $\mathcal{G}^c(\sigma^2)$ with $\Delta\sigma^2 = 0.01$ dB.

Integrity of choice of ΔR: We tentatively select resolution $\Delta R = 0.1$, and investigate the integrity of this ΔR over the 64 different settings above. After the coding rate sequence $\mathbf{R}^* = (R_1^*, \cdots, R_T^*)$ is obtained by DP for each setting, we randomly perturb R_t^* by $R_p(t) = R_t^* + \beta_t,\ t = 1, ..., T$, where $R_p(t)$ is the *perturbed coding rate*, the bias is $\beta_t \in \left[-\frac{\Delta R}{2}, +\frac{\Delta R}{2}\right]$, and $\mathbf{R}_p = (R_p(1), \cdots, R_p(T))$ is called the *perturbed coding rate sequence*. After randomly generating 100 different perturbed coding rate sequences $\mathbf{R}_p$, we calculate the aggregate coding rate (14), R_{agg}^p, of each $\mathbf{R}_p$; we only consider the perturbed coding rate sequences that achieve EMSE no greater than the optimal coding rate sequence $\mathbf{R}^*$ given by DP. The bottom panel of Fig. 7 plots the empirical PMF of ΔR_{agg}, where $\Delta R_{agg} = R_{agg}^p - R_{agg}^*$ and $R_{agg}^* = ||\mathbf{R}^*||_1$. Roughly 15% of cases in our simulation yield $\Delta R_{agg} < 0$ (meaning that the perturbed coding rate sequence has lower R_{agg}), while for the other 85% cases, $\mathbf{R}^*$ has lower R_{agg}. Considering the resolution $\Delta R = 0.1$, we can see that the perturbed sequences are only marginally better than $\mathbf{R}^*$. Hence, we verified the integrity of $\Delta R = 0.1$.

D. Proof of Lemma 1

Proof: We show that our DP scheme (17) fits into Bertsekas' formulation [27], which has been proved to be optimal. Under Bertsekas' formulation, our decision variable is the coding rate R_t and our state is the scalar channel noise variance σ_t^2. Our next-state function is the lossy SE (12) with the distortion D_t being calculated from the RD function given the decision variable R_t. Our additive cost associated with the dynamic system is $b \times \mathbb{1}_{R_t \neq 0} + R_t$. Our control law maps the state σ_t^2 to a decision (the coding rate R_t). Therefore, our DP formulation (17) fits into the optimal DP formulation of Bertsekas [27]. Hence, our DP formulation (17) is also optimal *for the discretized search spaces of R_t and σ_t^2.* ■

E. Proof of Theorem 1

Proof: Our proof is based on the assumption that lossy SE (12) holds. Consider the geometry of the SE incurred by

[8]Note that when calculating Φ^f, we are still using the corresponding interpolation scheme. Although this comparison is not ideal, we believe it still provides the reader with enough insight.

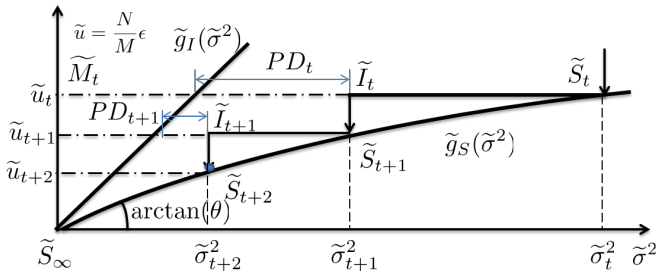

Fig. 8. Illustration of the evolution of u_t. The vertical axis shows $u_t = \frac{N}{M}\text{EMSE} = \frac{N}{M}\epsilon_t$. The solid lines with arrows denote the lossy SE associated with a coding rate sequence and dashed-dotted lines are auxiliary lines.

$\mathbf{R}^*$ for arbitrary iterations t and $t+1$, as shown in Fig. 8. Let $\widetilde{S}_t = (\widetilde{\sigma}_t^2, \widetilde{u}_t)$ and R_t^* be the state and the optimal coding rate at iteration t, respectively. We know that the slope of $\widetilde{g}_I(\cdot)$ is $\widetilde{g}_I'(\cdot) = 1$. Hence, the length of line segment $\widetilde{M}_t\widetilde{I}_t$ is $\widetilde{\sigma}_{t+1}^2 = \widetilde{u}_t + PD_t$. That is

$$PD_t = \widetilde{\sigma}_{t+1}^2 - \widetilde{u}_t. \tag{34}$$

Similarly, we obtain

$$PD_{t+1} = \widetilde{\sigma}_{t+2}^2 - \widetilde{u}_{t+1}, \tag{35}$$

where $\widetilde{u}_{t+1}$ and $\widetilde{\sigma}_{t+1}^2$ obey

$$\widetilde{\sigma}_{t+1}^2 = \widetilde{g}_S^{-1}(\widetilde{u}_{t+1}). \tag{36}$$

Recall that, according to Taylor's theorem (21), we obtain that

$$\widetilde{g}_S^{-1}(\widetilde{u}_{t+1}) = \frac{1}{\theta}\widetilde{u}_{t+1} + C\widetilde{u}_{t+1}^2, \tag{37}$$

with θ defined in (22). Although C depends on $\widetilde{u}_{t+1}$, it is uniformly bounded, i.e., $C \in [-B, B]$ for some $0 \le B < \infty$.

Fixing $\widetilde{u}_t = \frac{N}{M}\epsilon_t^*$ and $\widetilde{u}_{t+2} = \frac{N}{M}\epsilon_{t+2}^*$, we explore different distortions D_t and D_{t+1} that obey (34)–(36). According to Definition 2, among distortions that obey (34)–(36), the optimal D_t^* and D_{t+1}^* correspond to the smallest aggregate rate at iterations t and $t+1$, $R_t + R_{t+1}$. Considering (24), we have

$$R_t + R_{t+1} = \left(\frac{1}{2}\log_2\left(\frac{C_1}{D_t}\right) + \frac{1}{2}\log_2\left(\frac{C_1}{D_{t+1}}\right)\right)(1 + o_t(1)).$$

Therefore, in the large t limit, minimizing $R_t + R_{t+1}$ is identical to maximizing the product $D_t D_{t+1}$. Considering (34)–(36), our optimization problem becomes maximization over $F(\widetilde{u}_{t+1})$, where

$$F(\widetilde{u}_{t+1}) = (\widetilde{\sigma}_{t+2}^2 - \widetilde{u}_{t+1})(\widetilde{g}_S^{-1}(\widetilde{u}_{t+1}) - \widetilde{u}_t). \tag{38}$$

Invoking Taylor's theorem (37) and considering that $C \in [-B, B]$, we solve the optimization problem (38) in two extremes: one with $C = B$ and the other with $C = -B$.

In the case of $C = B$, we obtain

$$F(\widetilde{u}_{t+1}) = -\frac{1}{\theta}\widetilde{u}_{t+1}^2 + \frac{1}{\theta}\widetilde{u}_{t+1}\widetilde{\sigma}_{t+2}^2 + B\widetilde{\sigma}_{t+2}^2\widetilde{u}_{t+1}^2 - B\widetilde{u}_{t+1}^3 - \widetilde{u}_t\widetilde{\sigma}_{t+2}^2 + \widetilde{u}_t\widetilde{u}_{t+1}.$$

The maximum of $F(\widetilde{u}_{t+1})$ is achieved when $F'(\widetilde{u}_{t+1}) = 0$. That is,

$$F'(\widetilde{u}_{t+1}) = -3B\widetilde{u}_{t+1}^2 + \left(2B\widetilde{\sigma}_{t+2}^2 - \frac{2}{\theta}\right)\widetilde{u}_{t+1} + \frac{\widetilde{\sigma}_{t+2}^2}{\theta} + \widetilde{u}_t = 0. \tag{39}$$

Considering that $0 < \widetilde{u}_{t+1} < \widetilde{u}_t$, the root of the quadratic equation (39) is

$$\widetilde{u}_{t+1}^* = \frac{1}{3B}\left[\left(B\widetilde{\sigma}_{t+2}^2 - \frac{1}{\theta}\right) + A\right], \tag{40}$$

where

$$A = \sqrt{\left(B\widetilde{\sigma}_{t+2}^2 - \frac{1}{\theta}\right)^2 + 3B\left(\frac{\widetilde{\sigma}_{t+2}^2}{\theta} + \widetilde{u}_t\right)}. \tag{41}$$

We can further simplify (41) as

$$\begin{aligned} A =& \frac{1}{\theta}\sqrt{1 + B(\theta\widetilde{\sigma}_{t+2}^2 + B\theta^2\widetilde{\sigma}_{t+2}^4 + 3\theta^2\widetilde{u}_t)} \\ =& \frac{1}{\theta}\left[1 + \frac{B}{2}(\theta\widetilde{\sigma}_{t+2}^2 + B\theta^2\widetilde{\sigma}_{t+2}^4 + 3\theta^2\widetilde{u}_t)\right] + O(\widetilde{u}_t^2), \end{aligned} \tag{42}$$

Plugging (42) into (40),

$$\widetilde{u}_{t+1}^* = \frac{1}{2}(\widetilde{\sigma}_{t+2}^2 + \theta\widetilde{u}_t) + O(\widetilde{u}_t^2). \tag{43}$$

Plugging (43) into (34) and (35),

$$\begin{aligned} PD_t^* &= \frac{1}{2\theta}(\widetilde{\sigma}_{t+2}^2 - \widetilde{u}_t\theta) + O(\widetilde{u}_t^2), \\ PD_{t+1}^* &= \frac{1}{2}(\widetilde{\sigma}_{t+2}^2 - \widetilde{u}_t^2\theta) + O(\widetilde{u}_t^2), \end{aligned}$$

which leads to

$$\frac{D_{t+1}^*}{D_t^*} = \theta(1 + O(\widetilde{u}_t)). \tag{44}$$

These steps provided the optimal relation between D_t^* and D_{t+1}^* when $C = B$. For the other extreme case, $C = -B$, similar steps will lead to (44), where the differences between the results are higher order terms. Note that for any $C \in [-B, B]$ the higher order term is bounded between the two extremes. Hence, the optimal D_t^* and D_{t+1}^* follow (44) leading to the first part of the claim (25). Considering (24) and (44),

$$R_{t+1}^* - R_t^* = \frac{1}{2}\log_2\left(\frac{1}{\theta}\right)(1 + o_t(1)).$$

Therefore, we obtain the second part of the claim (26). ■

F. Proof of Theorem 2

Proof: Our proof is based on the assumption that lossy SE (12) holds. Let us focus on an optimal coding rate sequence $\mathbf{R}^* = (R_1^*, \cdots, R_T^*)$. Applying Taylor's theorem to calculate the ordinate of point $\widetilde{S}_{t+1}$ using its abscissa (Fig. 8), we obtain

$$\widetilde{u}_{t+1}^* = \theta(\widetilde{u}_t^* + PD_t^*) + O((\widetilde{u}_t^*)^2). \tag{45}$$

Therefore,

$$\frac{\widetilde{u}_{t+1}^*}{\widetilde{u}_t^*} = \theta + \frac{\theta PD_t^*}{\widetilde{u}_t^*} + O(\widetilde{u}_t^*). \tag{46}$$

Similarly, we obtain

$$\frac{\widetilde{u}_{t+2}^*}{\widetilde{u}_{t+1}^*} = \theta + \frac{\theta PD_{t+1}^*}{\widetilde{u}_{t+1}^*} + O(\widetilde{u}_t^*). \tag{47}$$

Plugging (44) and (45) into (47), we obtain

$$\begin{aligned}\frac{\widetilde{u}^*_{t+2}}{\widetilde{u}^*_{t+1}} &= \theta + \frac{\theta P D^*_t (1+O(u^*_t))}{\widetilde{u}^*_t + P D^*_t + O((\widetilde{u}^*_t)^2)} + O(\widetilde{u}^*_t) \\ &= \theta + \frac{\theta P D^*_t}{\widetilde{u}^*_t + P D^*_t} + O(\widetilde{u}^*_t).\end{aligned} \tag{48}$$

On the other hand, $\lim_{t\to\infty} \frac{\widetilde{u}^*_{t+1}}{\widetilde{u}^*_t} = \lim_{t\to\infty} \frac{\widetilde{u}^*_{t+2}}{\widetilde{u}^*_{t+1}}$. Therefore, considering (46) and (48), we obtain

$$\lim_{t\to\infty} \frac{\theta P D^*_t}{\widetilde{u}^*_t} = \lim_{t\to\infty} \frac{\theta P D^*_t}{\widetilde{u}^*_t + P D^*_t},$$

which leads to $\lim_{t\to} \frac{D^*_t}{\widetilde{u}^*_t} = 0$. We obtain (27) by noting that the optimal EMSE at iteration t is $\epsilon^*_t = \frac{M}{N}\widetilde{u}^*_t$. Plugging (27) into (46), we obtain (28). ∎

REFERENCES

[1] P. Han, J. Zhu, R. Niu, and D. Baron, "Multi-processor approximate message passing using lossy compression," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Process. (ICASSP)*, Shanghai, China, Mar. 2016, pp. 6240–6244.

[2] J. Zhu, A. Beirami, and D. Baron, "Performance trade-offs in multi-processor approximate message passing," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Barcelona, Spain, July 2016, pp. 680–684.

[3] D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[4] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.

[5] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing," *Proc. Nat. Academy Sci.*, vol. 106, no. 45, pp. 18914–18919, Nov. 2009.

[6] A. Montanari, "Graphical models concepts in compressed sensing," *Compressed Sensing: Theory and Applications*, pp. 394–438, 2012.

[7] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 57, no. 2, pp. 764–785, Feb. 2011.

[8] C. Rush and R. Venkataramanan, "Finite-sample analysis of approximate message passing," *Arxiv preprint arXiv:1606.01800*, June 2016.

[9] T. Tanaka, "A statistical-mechanics approach to large-system analysis of CDMA multiuser detectors," *IEEE Trans. Inf. Theory*, vol. 48, no. 11, pp. 2888–2910, Nov. 2002.

[10] D. Guo and S. Verdú, "Randomly spread CDMA: Asymptotics via statistical physics," *IEEE Trans. Inf. Theory*, vol. 51, no. 6, pp. 1983–2010, June 2005.

[11] D. Guo and C. C. Wang, "Multiuser detection of sparsely spread CDMA," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 3, pp. 421–431, Apr. 2008.

[12] D. Guo, D. Baron, and S. Shamai, "A single-letter characterization of optimal noisy compressed sensing," in *Proc. 47th Allerton Conf. Commun., Control, and Comput.*, Sept. 2009, pp. 52–59.

[13] S. Rangan, A. K. Fletcher, and V. K. Goyal, "Asymptotic analysis of MAP estimation via the replica method and applications to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 58, no. 3, pp. 1902–1923, Mar. 2012.

[14] J. Zhu and D. Baron, "Performance regions in compressed sensing from noisy measurements," in *Proc. IEEE Conf. Inf. Sci. Syst. (CISS)*, Baltimore, MD, Mar. 2013.

[15] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová, "Probabilistic reconstruction in compressed sensing: Algorithms, phase diagrams, and threshold achieving matrices," *J. Stat. Mech. – Theory E.*, vol. 2012, no. 08, pp. P08009, Aug. 2012.

[16] J. Mota, J. Xavier, and P. Aguiar, "Distributed basis pursuit," *IEEE Trans. Signal Process.*, vol. 60, no. 4, pp. 1942–1956, Apr. 2012.

[17] S. Patterson, Y. C. Eldar, and I. Keidar, "Distributed compressed sensing for static and time-varying networks," *IEEE Trans. Signal Process.*, vol. 62, no. 19, pp. 4931–4946, Oct. 2014.

[18] P. Han, R. Niu, M. Ren, and Y. C. Eldar, "Distributed approximate message passing for sparse signal recovery," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Atlanta, GA, Dec. 2014, pp. 497–501.

[19] C. Ravazzi, S. M. Fosson, and E. Magli, "Distributed iterative thresholding for ℓ_0/ℓ_1 -regularized linear inverse problems," *IEEE Trans. Inf. Theory*, vol. 61, no. 4, pp. 2081–2100, Apr. 2015.

[20] P. Han, R. Niu, and Y. C. Eldar, "Communication-efficient distributed IHT," in *Proc. Signal Process. with Adaptive Sparse Structured Representations Workshop (SPARS)*, Cambridge, United Kingdom, July 2015.

[21] G. J. Pottie and W. J. Kaiser, "Wireless integrated network sensors," *Commun. ACM*, vol. 43, no. 5, pp. 51–58, May 2000.

[22] D. Estrin, D. Culler, K. Pister, and G. Sukhatme, "Connecting the physical world with pervasive networks," *IEEE Pervasive Comput.*, vol. 1, no. 1, pp. 59–69, Jan. 2002.

[23] "Amazon EC2," https://aws.amazon.com/ec2/.

[24] T. Berger, *Rate Distortion Theory: Mathematical Basis for Data Compression*, Prentice-Hall Englewood Cliffs, NJ, 1971.

[25] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, New York, NY, USA: Wiley-Interscience, 2006.

[26] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*, Kluwer, 1993.

[27] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, vol. 1, Athena Scientific Belmont, MA, 1995.

[28] Y. Ma, D. Baron, and D. Needell, "Two-part reconstruction with noisy-sudocodes," *IEEE Trans. Signal Process.*, vol. 62, no. 23, pp. 6323–6334, Dec. 2014.

[29] I. Das and J. E. Dennis, "Normal-boundary intersection: A new method for generating the Pareto surface in nonlinear multicriteria optimization problems," *SIAM J. Optimization*, vol. 8, no. 3, pp. 631–657, Aug. 1998.

[30] P. Frasca, R. Carli, F. Fagnani, and S. Zampieri, "Average consensus on networks with quantized communication," *Int. J. Robust Nonlinear Control*, vol. 19, no. 16, pp. 1787–1816, Nov. 2008.

[31] D. Thanou, E. Kokiopoulou, Y. Pu, and P. Frossard, "Distributed average consensus with quantization refinement," *IEEE Trans. Signal Process.*, vol. 61, no. 1, pp. 194–205, Jan. 2013.

[32] J. Tan, D. Carmon, and D. Baron, "Signal estimation with additive error metrics in compressed sensing," *IEEE Trans. Inf. Theory*, vol. 60, no. 1, pp. 150–158, Jan. 2014.

[33] Y. Ma, J. Zhu, and D. Baron, "Approximate message passing algorithm with universal denoising and Gaussian mixture learning," *IEEE Trans. Signal Process.*, vol. 65, no. 21, pp. 5611–5622, Nov. 2016.

[34] Y. Ma, D. Baron, and A. Beirami, "Mismatched estimation in large linear systems," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, July 2015, pp. 760–764.

[35] C. Weidmann and M. Vetterli, "Rate distortion behavior of sparse sources," *IEEE Trans. Inf. Theory*, vol. 58, no. 8, pp. 4969–4992, Aug. 2012.

[36] S. Arimoto, "An algorithm for computing the capacity of an arbitrary discrete memoryless channel," *IEEE Trans. Inf. Theory*, vol. 18, no. 1, pp. 14–20, Jan. 1972.

[37] R. E. Blahut, "Computation of channel capacity and rate-distortion functions," *IEEE Trans. Inf. Theory*, vol. 18, no. 4, pp. 460–473, July 1972.

[38] K. Rose, "A mapping approach to rate-distortion computation and analysis," *IEEE Trans. Inf. Theory*, vol. 40, no. 6, pp. 1939–1952, Nov. 1994.

[39] Y. Linde, A. Buzo, and R. M. Gray, "An algorithm for vector quantizer design," *IEEE Trans. Commun.*, vol. 28, no. 1, pp. 84–95, Jan. 1980.

[40] R. M. Gray, "Vector quantization," *IEEE ASSP Mag.*, vol. 1, no. 2, pp. 4–29, Apr. 1984.

[41] B. Widrow and I. Kollár, *Quantization Noise: Roundoff Error in Digital Computation, Signal Processing, Control, and Communications*, Cambridge University press, 2008.

[42] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Trans. Inf. Theory*, vol. IT-44, pp. 2325–2383, Oct. 1998.

[43] Y. Wu and S. Verdú, "MMSE dimension," *IEEE Trans. Inf. Theory*, vol. 57, no. 8, pp. 4857–4879, Aug. 2011.

[44] Texas Instruments, *A true system-on-chip solution for 2.4-GHz IEEE 802.15.4 and ZigBee applications*, Apr. 2009, Rev. B.

[45] "Pearson product-moment correlation coefficient," https://en.wikipedia.org/wiki/Pearson_product-moment_correlation_coefficient.